

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23

EMCR:034-2
EMC-97-159 CON2

MAPPING OF HOSTS TO LOGICAL STORAGE UNITS AND
DATA STORAGE PORTS IN A DATA PROCESSING SYSTEM

BY:

Steven M. Blumenau

and

Yoav Raz

EXPRESS MAIL NO. EV 318623724 US

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24

RELATED APPLICATIONS

This application is a continuation of application Ser. 10/139,490 filed May 6, 2002, which is a continuation of Ser. No. 09/106,299 filed June 29, 1998, issued as U.S. Patent 6,421,711, incorporated herein by reference.

BACKGROUND OF THE INVENTION

Technical Field.

The present invention relates generally to data processing networks and data storage subsystems, and more particularly to a data processing network in which a large number of hosts can access volumes of data storage in a data storage subsystem.

Description of the Related Art.

Due to advances in computer technology, there has been an ever increasing need for data storage in data processing networks. In a typical data processing network, there has been an increase in the number of volumes of data storage and an increase in the number of hosts needing access to the volumes. This has been especially true for networks of workstations. Not only have a greater number of workstations been added to the typical network, but also the increase in data processing capabilities of a typical workstation has required more data storage per workstation for enhanced graphics and video applications.

The increased demand for data storage in a network is typically met by using more storage servers in the network or by using storage servers of increased storage capacity and data transmission bandwidth. From the standpoint of cost of storage, either

1 of these solutions appears to be satisfactory. However, a greater number of storage
2 servers in a network substantially increases the cost of managing the storage. This
3 increased cost of management often appears some time after installation, when one of the
4 servers reaches its capacity and some of its volumes must be reassigned to less heavily
5 loaded servers. Network administrators aware of the cost of storage management realize
6 that network storage should be consolidated to the minimum possible number of servers.
7 The management problem is reduced by reducing the number of objects to be managed.

8
9 Due to the storage needs of present networks and the desire to consolidate servers,
10 it is practical to provide a single storage subsystem with up to 20 terabytes (TB) storage,
11 or approximately 4000 logical volumes. It may be possible for any host to have access to
12 any volume in a data storage subsystem to which the host has access. However, it may be
13 desirable to restrict the set of volumes that can be seen by any one host. Restricted access
14 is desirable for security of private data. For example, private volumes should be assigned
15 to each host for storage of private data, and other hosts should not be permitted to see or
16 modify the private volumes of other hosts. Moreover, the "boot" process for a host is
17 slowed down by searching for and reporting all the volumes to which the host has access.
18 Certain operating systems are limited by the number of storage devices that they can
19 manage at a given period of time, and for a host running such an operating system, it is
20 not only desirable but also necessary to limit the number of volumes that the host can
21 access.

22

1 It is possible to restrict access of a host to a limited set of logical volumes in the
2 data storage subsystem by restricting the set of logical volumes accessible through a
3 particular port adapter of the storage subsystem and linking the host to only that
4 particular port adapter. For convenience, however, there should not be any restrictions on
5 which logical storage volumes are accessible from each port adapter. Otherwise, during a
6 reconfiguration of the data processing system, it may be necessary to physically switch
7 the links that are connected to the network ports of the hosts or the port adapters, for
8 example by manually disconnecting and reconnecting the links to the ports. Even in the
9 case where the data network has a fabric for automatically establishing a link between
10 any of the hosts and any of the port adapters, the physical possibility of any port adapter
11 to access any logical storage volume provides alternative data paths that could be used in
12 case of port adapter failure or port adapter congestion. For example, if a host sends a data
13 access request to a port adapter and receives a busy response from the port adapter, then
14 the host can send the data access request to another port adapter. Port adapter congestion
15 is likely, for example, if the storage subsystem is a continuous media server, in which
16 video data is often streamed through a single port adapter to a host for a relatively long
17 period of time.

18
19 In open network systems, it is known to use authentication and authorization
20 protocols in order to authenticate that a request for access to a specified file originates
21 from a particular host, and once the request for access is authenticated, to check whether
22 the host is authorized to access the specified file. For example, a network server

1 authenticates the request by checking whether a password in the request matches the
2 hosts' password stored in a client directory, and the network server authorizes the request
3 by checking a file directory to determine whether the host is listed in the file directory as
4 having access rights to the specified file. However, the use of high-level authentication
5 and authorization procedures for discriminating among all access requests by the hosts to
6 the logical storage volumes would unduly burden the host and the storage subsystem.
7 What is desired is a method that may be transparent to any high-level file system
8 procedures that may be used by the hosts for managing access to files stored in the logical
9 volumes to which a host is permitted to access. The method should restrict the logical
10 storage volumes seen by the host during a boot operation, and seen by the operating
11 system when the operating system determines what logical volumes are accessible to the
12 host.

14 SUMMARY OF THE INVENTION

15 In accordance with one aspect of the invention, there is provided a data storage
16 subsystem that includes data storage and a storage controller coupled to the data storage
17 for controlling access to the data storage. The storage controller has at least one physical
18 data port for connecting the storage controller into a data network for data transmission
19 between the data storage and host processors in the data network. The storage controller
20 is programmed to provide a plurality of virtual ports that are not physical ports in the data
21 network but that appear to the host processors to be physical ports in the data network
22 that provide access to the data storage and that are connected to the physical data port by

1 a switch in the storage controller for routing storage access requests from the physical
2 data port to the virtual ports.

3 In accordance with another aspect, the invention provides a data storage
4 subsystem including data storage and a storage controller coupled to the data storage for
5 controlling access to the data storage. The storage controller has at least one physical
6 data port for connecting the storage controller into a data network for data transmission
7 between the data storage and host processors in the data network. The storage controller
8 is programmed to provide a plurality of virtual ports that are not physical ports in the data
9 network but that appear to the host processors to be physical ports in the data network
10 that provide access to the data storage and that are connected to the physical data port by
11 a switch in the storage controller for routing storage access requests from the physical
12 data port to the virtual ports. The storage controller is programmed with a permanent
13 network name for each of the virtual ports, and with a temporary network address for
14 each of the virtual ports. The storage controller is programmed to route from the physical
15 data port to a specified virtual port a storage access request containing a temporary
16 address that specifies the specified virtual port. The storage controller is also
17 programmed to provide access at each virtual port to only a respective assigned subset of
18 the data storage. The storage controller is programmed to permit assignment of more
19 than one virtual port to each host processor such that each host processor may access
20 storage from every virtual port assigned to the host processor. The storage controller is
21 further programmed so that none of the virtual ports is assigned to more than one host
22 processor so that not more than one host processor may access the data storage from any

1 one of the virtual ports.

2 In accordance with another aspect, the invention provides a machine-readable
3 program storage device containing a program that is executable by a storage controller for
4 controlling access to data storage. The storage controller has at least one physical data
5 port for connecting the storage controller into a data network for data transmission
6 between the data storage and host processors in the data network. The program is
7 executable by the storage controller to provide a plurality of virtual ports that are not
8 physical ports in the data network but that appear to the host processors to be physical
9 ports in the data network that provide access to the data storage and that are connected to
10 the physical data port by a switch in the storage controller for routing storage access
11 requests from the physical data port to the virtual ports.

12 In accordance with yet another aspect, the invention provides a machine-readable
13 program storage device that is executable by a storage controller for controlling access to
14 data storage. The storage controller has at least one physical data port for connecting the
15 storage controller into a data network for data transmission between the data storage and
16 host processors in the data network. The program is executable by the storage controller
17 to provide a plurality of virtual ports that are not physical ports in the data network but
18 that appear to the host processors to be physical ports in the data network that provide
19 access to the data storage and that are connected to the physical data port by a switch in
20 the storage controller for routing storage access requests from the physical data port to the
21 virtual ports. The program is executable by the storage controller so that the storage
22 controller will have a permanent network name for each of the virtual ports and a

1 temporary network address for each of the virtual ports. The program is executable by
2 the storage controller to route from the physical data port to a specified virtual port a
3 storage access request containing a temporary address that specifies the specified virtual
4 port. The program is executable by the storage controller to provide access at each virtual
5 port to only a respective assigned subset of the data storage. The program is executable
6 by the storage controller to permit assignment of more than one virtual port to each host
7 processor such that each host processor may access storage from every virtual port
8 assigned to the host processor. The program is also executable by the storage controller
9 so that none of the virtual ports is assigned to more than one host processor so that not
10 more than one host processor may access the data storage from any one of the virtual
11 ports.

12 In accordance with still another aspect, the invention provides a method of
13 operating a storage controller for controlling access to data storage. The storage
14 controller has at least one physical data port for connecting the storage controller into a
15 data network for data transmission between the data storage and host processors in the
16 data network. The method includes the storage controller receiving storage access
17 requests from the host processors at the physical data port, inspecting network addresses
18 in the storage access requests to find network addresses of virtual ports in the storage
19 controller to which the storage access requests are directed, and controlling access to the
20 data storage in accordance with the network addresses of the virtual ports to which the
21 storage access requests are directed. The virtual ports are not physical data ports in the
22 data network, but the storage controller is operated to cause the virtual ports to appear to

1 the host processors to be physical ports in the data network that provide access to the data
2 storage and that are connected to the physical data port by a switch in the storage
3 controller for routing storage access requests from the physical data port to the virtual
4 ports.

5 In accordance with a final aspect, the invention provides a method of operating a
6 storage controller for controlling access to data storage. The storage controller has at
7 least one physical data port for connecting the storage controller into a data network for
8 data transmission between the data storage and host processors in the data network. The
9 method includes the storage controller receiving storage access requests from the host
10 processors at the physical data port, and inspecting network addresses in the storage
11 access requests to find network addresses of virtual ports in the storage controller to
12 which the storage access requests are directed, and controlling access to the data storage
13 in accordance with the network addresses of the virtual ports to which the storage access
14 requests are directed. The virtual ports are not physical data ports in the data network,
15 but the storage controller is operated to cause the virtual ports to appear to the host
16 processors to be physical ports in the data network that provide access to the data storage
17 and that are connected to the physical data port by a switch in the storage controller for
18 routing storage access requests from the physical data port to the virtual ports. The
19 storage controller maintains a permanent network name for each of the virtual ports and a
20 temporary network address for each of the virtual ports. The storage controller provides
21 access at each virtual port to only a respective assigned subset of the data storage. The
22 storage controller permits assignment of more than one virtual port to each host processor

1 such that each host processor may access storage from every virtual port assigned to the
2 host processor. Moreover, none of the virtual ports is assigned to more than one host
3 processor so that not more than one host processor may access the data storage from any
4 one of the virtual ports.

5
6

7 BRIEF DESCRIPTION OF THE DRAWINGS

8

9 Other objects and advantages of the invention will become apparent upon reading
10 the following detailed description with reference to the accompanying drawings wherein:

11

12 FIG. 1 is a block diagram of a data processing system including a cached storage
13 subsystem linked by a data network to a multiplicity of host processors;

14

15 FIG. 2 is a block diagram of the data processing system of FIG. 1 further showing
16 a typical fault-tolerant implementation for the data network;

17

18 FIG. 3 is a generalized block diagram of a node linked to the data network of FIG.

19 1;

20

21 FIG. 4 is a block diagram of the data processing system of FIG. 1 showing a port
22 adapter in the cached storage subsystem using a volume access table for assignment of
23 logical storage volumes to respective host controller ports;

1

2 FIG. 5 is a diagram showing a simplified construction that could be used for the
3 volume access table introduced in FIG. 4;

4

5 FIG. 6 is a diagram of a volume list entry defining a vector of logical storage
6 volumes in a disk spread;

7

8 FIG. 7 is a block diagram of the cached storage subsystem showing further data
9 structures used with the volume access table stored in port adapter memory;

10

11 FIG. 8 is a diagram showing a preferred construction for a volume access table
12 stored in the cache memory as introduced in FIG. 7;

13

14 FIG. 9 is a diagram showing a preferred construction for the volume access table
15 stored in port adapter memory;

16

17 FIG. 10 is a diagram of a directory for searching the volume table of FIG. 9;

18

19 FIG. 11 is a flowchart of a microcode routine executed by a port adapter of the
20 cached storage subsystem of FIG. 1 when using the volume access table of FIG. 5 or FIG.
21 9 in response to a "Report LUNs" or access request by a host processor;

22

1 FIG. 12 is a flowchart of a microcode routine executed by a port adapter of the
2 cached storage subsystem of FIG. 1 when converting the vector representation of FIG. 6
3 to a list of logical storage volume numbers;

4
5 FIG. 13 is a flowchart of a microcode routine executed by a port adapter of the
6 cached storage subsystem of FIG. 1 when determining whether a specified logical storage
7 volume is included in a disk spread defined by the vector representation of FIG. 6;

8
9 FIG. 14 is a graphical display of disk spreads in a two-dimensional volume space;

10
11 FIG. 15 is a graphical display of disk spreads in a three-dimensional volume
12 space;

13
14 FIG. 16 is a flowchart of a microcode routine executed by a port adapter of the
15 cached storage subsystem of FIG. 1 for volume configuration when installing the storage
16 subsystem into the data processing system of FIG. 1;

17
18 FIG. 17 is a flowchart of a microcode routine executed by a port adapter of the
19 cached storage subsystem of FIG. 1 when notified of a network state change such as a
20 host controller boot;

21
22 FIG. 18 is a flowchart of a procedure performed by a system administrator during

1 host controller replacement;

2

3 FIG. 19 is a diagram of a volume list that comprises a mapping table for mapping
4 logical unit numbers (LUNs) to logical storage volume numbers;

5

6 FIG. 20 is a diagram of a volume list that comprises a mapping table for mapping
7 ranges of LUNs to vectors of logical storage volume numbers;

8

9 FIG. 21 is a schematic diagram of the data processing system of FIG. 1 showing
10 one of the port adapters programmed to provide virtual ports for access to respective
11 groups of logical storage volumes;

12

13 FIG. 22 is a schematic diagram of a data processing system in which port adapters
14 are programmed to permit a host to access the same virtual ports through the physical
15 ports of more than one of the port adapters.

16

17 FIG. 23 is block diagram of data structures of volume access and mapping
18 information stored in memory of port adapters of the storage subsystem of FIG. 21 or
19 FIG. 22;

20

21 FIG. 24 is a diagram of a virtual port host table included in the volume access and
22 mapping information of FIG. 23;

1

2 FIG. 25 is a diagram of a virtual port mapping table included in the volume access
3 and mapping information of FIG. 23;

4

5 FIG. 26 is a flow chart of a routine performed by a port adapter when reporting
6 the virtual ports accessible to a host controller port;

7

8 FIG. 27 is a first sheet of a flow chart of a routine performed by a port adapter
9 when responding to a volume access request or “Report LUNs” request from a host;

10

11 FIG. 28 is a second sheet of the flow chart begun in FIG. 27;

12

13 FIG. 29 is a diagram of a display generated by a graphical user interface to show a
14 mapping relationship between logical volumes, storage adapter ports, and LUNs at a
15 virtual port;

16

17 FIG. 30 is a diagram of a display generated by a graphical user interface to
18 establish a mapping between volumes accessible at a virtual port for a host and the LUNs
19 at the virtual port;

20

21 FIG. 31 is a block diagram of a semiconductor integrated circuit chip capable of
22 authenticating itself in a secure fashion;

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

FIG. 32 is a block diagram showing semiconductor integrated circuit chips in accordance with FIG. 31 being used in the data processing system of FIG. 1 to permit a port adapter to authenticate the identity of host controllers;

FIG. 33 is a flowchart of a “challenge-response” procedure used by the port adapter and host controllers of FIG. 32 to permit the port adapter to authenticate the identity of a host controller;

FIG. 34 is a block diagram showing components of a Fibre Channel Frame;

FIG. 35 is a first sheet of a flowchart of a procedure used by the port adapter and host controllers of FIG. 32 to permit the port adapter to authenticate each message from a host controller;

FIG. 36 is a second sheet of the flowchart begun in FIG. 35;

FIG. 37 is a block diagram of a storage network configuration especially suited for up to 800 workstations;

FIG. 38 is a block diagram of a storage network configuration especially suited for up to 3300 workstations;

1
2 FIG. 39 is a block diagram of a storage area network providing fault tolerance and
3 failover capability for hosts and storage subsystems in the network; and
4

5 FIG. 40 is a block diagram of a storage area network similar to that shown in FIG.
6 39 but including additional network loops for higher throughput.
7

8 While the invention is susceptible to various modifications and alternative forms,
9 specific embodiments thereof have been shown in the drawings and will be described in
10 detail. It should be understood, however, that it is not intended to limit the invention to
11 the particular forms shown, but on the contrary, the intention is to cover all modifications,
12 equivalents, and alternatives falling within the scope of the invention as defined by the
13 appended claims.
14

15 DESCRIPTION OF THE PREFERRED EMBODIMENTS 16

17 With reference to FIG. 1 of the drawings, there is shown a cached storage
18 subsystem 20 connected via a data network 21 to a plurality of hosts 22, 23, 24, 25. The
19 cached storage subsystem 20 includes storage volumes 26 and a storage controller 27 for
20 controlling access of the hosts to the storage volumes. The storage volumes are logical
21 units of storage distributed over one more storage devices 28, 29, 30, and 31. The storage
22 devices are magnetic disk drives, optical disk drives, tape drives, solid-state memory

1 devices, or other storage devices capable of providing nonvolatile data storage. Presently
2 the preferred storage devices are magnetic disk drives each having a storage capacity of at
3 least 46 gigabytes.

4
5 The storage controller 27 includes a dual port cache memory 32, a plurality of
6 port adapters 35, 36, and a plurality of storage adapters 37, 38. The cache memory 32 is
7 accessed via any one of two back-plane busses 33, 34. Each of the port adapters 35, 36
8 link the data network 21 to each of the two back-plane busses 33, 34. Each of the storage
9 adapters 37, 38 links a respective set of the storage devices 28, 29, 30, 31 to each of the
10 two back-plane busses 33, 34. For example, the cached storage subsystem includes up to
11 eight storage adapters and up to eight port adapters, and each port adapter provides two
12 independent data ports to the data network.

13
14 When a port adapter 35 or 36 receives a storage access request from one of the
15 hosts 22, 23, 24, 25, the port adapter accesses a directory in the cache memory 32 to
16 determine whether or not the data to be accessed resides in the cache memory. If the data
17 to be accessed resides in the cache memory, then the port adapter accesses the data in the
18 cache memory. If the data to be accessed does not reside in the cache memory, then the
19 port adapter forwards a storage access request to the storage adapters 37, 38. One of the
20 storage adapters 37, 38 responds to the storage access request by performing a logical-to-
21 physical translation to determine where the data to be accessed resides on the storage
22 devices, and reads the data from the storage devices and writes the data to the cache

1 memory, for access by the port adapter. The storage adapters 37, 38 also perform a write-
2 back operation to ensure that that data written to the cache memory 32 by the port
3 adapters eventually becomes written to the storage volumes 26.

4
5 The cache memory 32 ensures that data frequently accessed by the hosts is likely
6 to be found in cache in order to avoid the data access time of the storage devices and in
7 order to minimize loading on the storage adapters and the port adapters. Consolidation of
8 network storage into a large cached storage subsystem provides a benefit that cache
9 resources are consolidated into one large cache, which is more efficient than a number of
10 smaller caches having in total the same cache memory capacity. A large cache is more
11 likely to contain the most recently accessed data than the combined cache memory of the
12 smaller caches.

13
14 The storage subsystem 20 is constructed for high data availability so that a single
15 high-capacity storage subsystem is at least as fault-tolerant as a local collection of
16 conventional network storage servers. Fault tolerance is ensured by dual, redundant
17 components and busses in the path from any one of the port adapters 35, 36 to any one of
18 the storage devices 28, 29, 30, and 31. Mirroring or RAID (redundant array of
19 inexpensive disks) techniques ensure that the storage adapters 37, 38 can recover data in
20 the event of failure of any one of the storage devices. In a similar fashion, the data
21 network 21 can be made fault tolerant by ensuring that each of the hosts 22, 23, 24, 25
22 has independent paths through the data network 21 to each of two of the port adapters 35,

1 36, as will be further described below with reference to FIG. 2.

2

3 In a preferred form of construction, the cache memory 32 is composed of dynamic
4 RAM memory cards mounted in a card-cage or main-frame, and the port adapters and
5 storage adapters are programmed micro-processor cards that are also mounted in the card-
6 cage or main-frame. Each port adapter 35, 36 has one or more processors for handling
7 the communication protocol of the data network 21 and communicating with the cache
8 memory busses 33, 34. Each storage adapter 37, 38 has one or more processors for
9 handling the communication protocol of the storage devices and for communicating with
10 the cache memory busses 33, 34. For example, the links between the storage adapters 37
11 and the storage devices 28, 29, 30, and 31 are FWD (fast, wide, differential) SCSI or
12 Fibre Channel fiber-optic loops. The port adapters 35, 36 can be programmed to
13 communicate with the network via any number of communication and/or network
14 protocols, such as Bus and Tag CKD, ESCON, SCSI, Ethernet, FDDI, ATM, DS1, DS3,
15 T3, TCP, UDP, NFS, SNMP, and Fibre Channel. Further details regarding the preferred
16 construction and operation of the cached storage subsystem 20 are disclosed in Yanai et
17 al., U.S. Patent 5,206,939, issued April 27, 1993; Yanai et al. U.S. Patent No. 5,335,352,
18 issued Aug. 2, 1994; and Yanai et al. U.S. Patent 5,381,539, issued Jan. 10, 1995; all
19 incorporated herein by reference.

20

21 Referring to FIG. 2, there is shown a fault-tolerant way of connecting a large
22 number of hosts 22, 23, 24, 25 to the cached storage subsystem 20. In this example, the

1 data network 21 includes a respective loop 41, 42, 43, 44 connected to each port of the
2 port adapters 35, 36, and each host has at least two ports, each of which is connected to a
3 respective loop connected to the port of a different one of the port adapters. Therefore, if
4 there is a single failure of any one of the loops or a single failure of any one of the port
5 adapters, there will still be an operational path from each host to the internal back-plane
6 busses (33, 34 in FIG. 1) in the cached disk storage subsystem. The loops 41, 42, 43, 44,
7 for example, operate in accordance with the Ethernet or Fibre Channel standards. Each
8 loop could connect up to fifty hosts, for a total of 400 hosts connected via dual-redundant
9 paths from eight port adapters. For workstation hosts where dual-redundant paths would
10 not be needed, only one port of each workstation could be connected to the network 21,
11 so that up to 800 workstations could be connected to the storage subsystem by connecting
12 this single port of each workstation to only one of sixteen loops, as will be further
13 described below with reference to FIG. 37.

14
15 It is possible to replace each of the loops 41-44 in FIG. 2 with a switch, or to use
16 switches together with loops for connecting hosts to the storage subsystem, as will be
17 further described below with reference to FIG. 38. In general, for a given number of
18 ports, a loop will be less expensive than a switch, but a switch may provide more
19 bandwidth than a loop. The additional bandwidth of a switch may be needed for ensuring
20 concurrent host access to the storage subsystem or for supporting bandwidth-intensive
21 applications such as interactive video applications.

1 In general, the data network 21 of FIG. 1 could have various topologies and
2 simultaneously use a number of different communication protocols. For example, the
3 data network 31 could connect ports of various devices by dedicated port-to-port
4 connections (i.e., so-called point-to-point connections), loops, or switches, or
5 combinations of these connections. In addition to switches, the data network 21 could
6 include other connectivity devices such as hubs, bridges and routers. The data network
7 31 could include additional processors or computers for buffering or streaming data from
8 the cached storage subsystem to hosts or to archival storage such as a tape library. For
9 example, the use of a cached storage subsystem in connection with stream servers, a tape
10 library, and an ATM switch, for continuous-media isochronous data access, and on-line
11 backup storage, is described in Vishlitzky et al., U.S. Patent 5,737,747 issued April 7,
12 1998.

13
14 The data network 21 could also include one or more additional cached storage
15 subsystems, preferably at different geographical locations, to provide redundancy and
16 protection from a catastrophic failure of the cached storage subsystem 20, as will be
17 further described below with reference to FIGs. 39 and 40. The cached storage
18 subsystems could be linked for automatic remote mirroring of data to provide recovery
19 from a disaster, as described in Yanai et al., U.S. Patent 5,544,347 issued Aug. 6, 1996,
20 incorporated herein by reference, and in Yanai et al., U.S. Patent 5,742,792 issued April
21 21, 1998 (Ser. No. 08/654,511 filed May 28, 1996), incorporated herein by reference.

1 Because the data network 21 could have various topologies, it is desirable to
2 provide a facility for enabling any device in the network to determine the configuration of
3 at least that portion of the network that is accessible to the device. With reference to FIG.
4 3, for example, there is shown a node 50 having a number of ports, including ports 51 and
5 52, which may be connected to the data network 21. In this context, the node 50 can
6 represent any device having ports which may be connected to the data network 21. The
7 node also has a number of entities 53, 54, 55, 56, which are accessible through one or
8 more of the ports, and which are identified by logical unit numbers (LUNs), each of
9 which is unique for any single node. For example, if the node 50 is a storage subsystem,
10 each logical storage volume is assigned a unique LUN. As will be further described
11 below, it is also possible to establish a mapping between the LUNs as specified by or
12 reported to a host and the logical volumes, so that one logical volume could be mapped to
13 more than one LUN. Each port 51, 52 is constructed or programmed to respond to a
14 “Report LUNs” command by returning a list of LUNs which are accessible from the port.
15 In the fashion, a host can send “Report LUNs” commands to each port of each storage
16 subsystem to which it is connected, to obtain a list of the logical volumes to which the
17 host is connected. This is typically done by a host operating system at “boot” time.

18

19 The format for the Report LUNs command and the manner of its use by the host
20 is determined by the protocol used by the network. For example, for a host connected to
21 a storage node via a SCSI link, then the host addresses the storage volume by a
22 Bus/Target/LUN triple, and the host can send a SCSI “Report LUNs” command over the

1 bus to each target on the bus to obtain a list of LUNs accessible through each target.
2 Such a Report LUNs command is subsumed in the Fibre Channel standards, which define
3 additional capabilities for permitting a host to determine the present configuration of a
4 network to which the host is connected. The Fibre Channel standards are currently being
5 developed by the American National Standards Institute (ANSI).

6

7 In a Fibre Channel network, each port has a 64-bit port identifier called a “World
8 Wide Name” (WWN). To ensure that the WWN is unique, for example, it may include a
9 unique identifier of the manufacturer of the device including the port (e.g., an
10 Organizationally Unique Identifier as registered with the IEEE in New York, N.Y.), and
11 the manufacturer’s serial number of the device. A logical storage volume in a Fibre
12 Channel storage subsystem therefore can be identified by the combination of the WWN
13 of the storage subsystem and the LUN of the logical volume.

14

15 In a Fiber Channel network, the portion of the network connected to each port has
16 one of four distinct topologies; namely, point-to-point, private loop, public loop, and
17 fabric. A point-to-point topology has one port connected directly to one other port. A
18 loop is a single simplex media linking three or more nodes such that transmission can
19 occur between a single pair of nodes on the media at any given time. A hub is a specific
20 implementation of a loop in which a node can be inserted into or removed from the hub
21 without breaking the loop. A private loop is a dedicated or closed loop. A public loop is
22 a loop attached to a node of a fabric. A fabric is a topology that permits multiple

1 simultaneous transmissions between pairs of nodes connected to it. The fabric routes or
2 switches data frames from a source node to a destination node. A switch is an
3 implementation of the fabric topology.

4
5 In a Fiber Channel network, a node of a host directly connected to a loop can
6 determine all other nodes in the loop by polling. A response from a node on the loop can
7 indicate whether or not the responding node is or is not a fabric port. When a node is
8 identified, its associated LUNs can be requested by sending a Report LUNs request to the
9 node. A node of a fabric port can be further interrogated about the identity of other nodes
10 of the fabric to which the node of the fabric port is directly connected through the fabric.
11 For example, associated with the fabric is a “name server” that will respond with a list of
12 nodes directly accessible through the fabric. The name server has a predefined address on
13 the fabric. A node of the host can also probe the configuration of a point-to-point
14 connection or loop connected to each node directly accessible through the fabric, by
15 addressing a port of the fabric to route interrogation requests to other ports.

16
17 A request from one port to another on a loop or fabric must identify the
18 destination port so that only the destination port will respond to the request. A request
19 from one port to another on a loop or fabric must also identify the source of the request so
20 that a response can be directed back to the source. Although the WWNs of the source
21 and destination ports could be used in each request to uniquely identify the source and
22 destination ports, each WWN contains so many bits that an undue amount of transmission

1 bandwidth and data processing capability would be expended if the source and
2 destination WWNs were used in each request. Therefore, it is desirable to assign a
3 temporary identifier to each port in the data network in such a fashion that the identifier is
4 unique to the configuration of the network at any given time, but not necessarily unique
5 to each port for all time. Therefore, such a temporary identifier can have fewer bits than a
6 WWN yet uniquely identify a source or desired destination of a request transmitted
7 through the network.

8
9 The Fiber Channel standards specify that data is transmitted over the Fiber
10 Channel network in packets having a predefined format, and such a packet is called a
11 "frame." Each frame includes a source address (S_ID) and a destination address (D_ID).
12 The S_ID is a temporary identifier of the port which is the source of the frame, and the
13 D_ID is a temporary identifier of the port which is the desired destination of the frame.
14 Further details regarding the format of the frame are described below with reference to
15 FIG. 34.

16
17 The use of temporary rather than permanent identifiers for source and destination
18 addresses in the network introduces the problem of assigning and reassigning the
19 temporary identifiers when the configuration of the network is defined and changed. For
20 example, when a private loop is changed into a public loop by connecting a port of a
21 switch to the loop, temporary addresses must be assigned to other ports of the switch and
22 to the ports of devices that become connected to those other ports.

1 If a Fiber Channel network includes a fabric, then the temporary identifiers of ports
2 in the network are assigned by the fabric at the time of a fabric login if the fabric supports
3 a fabric login. Otherwise, the temporary identifiers must be known by the ports implicitly.
4 For example, in a fabric containing only one switch, when that one switch is powered on
5 or reset, then that one switch implicitly knows its own ports, and a “domain controller” of
6 the switch searches for nodes connected to the switch ports, assigns temporary IDs to all
7 the nodes that become known to it, and reports to the ports that it has found the
8 temporary IDs assigned to them. During this fabric login process, a “name server” of the
9 fabric may build a table of the WWNs of the ports that are known to the fabric and their
10 corresponding temporary IDs. If the Fibre Channel network includes multiple switches,
11 then only one of them, designated as the master switch, assigns ranges of temporary IDs
12 to be used by each of the switches during the fabric login process. The Fibre Channel
13 network can be constructed so that the name server automatically obtains information
14 from the domain controller when the domain controller learns about the ports that become
15 connected to the network, or the Fibre Channel network can be constructed so that a port
16 must log into the domain controller to obtain its temporary ID, and then log into the name
17 server so that the name server obtains the temporary ID of the port. The first case is
18 known as an implicit name server login, and the second case is known as an explicit name
19 server login.

20

21 After an initial fabric login process, it is possible for the configuration of the
22 network to become changed as a result of addition or replacement of devices having

1 ports connected to the network. The data processing system could be shut down and then
2 reconfigured so that the new ports become known implicitly, or a fabric login could be
3 initiated so that temporary IDs become assigned to the new ports. In the case of a data
4 network having a fabric supporting a fabric login, it is also possible for a device added to
5 the network to initiate a process of obtaining temporary IDs for its ports by requesting a
6 “port login” from the fabric.

7
8 The Fiber Channel specifications provide a mechanism for the network to
9 automatically detect certain changes of state which may indicate that the configuration of
10 the system has changed. For example, idle signals are transmitted over the links to enable
11 detection of link failure. Frame transmission errors are detected by cyclic redundancy
12 checks and sequence identification numbers in the frames. When transmission over a link
13 is restored after detection of a link failure, a fabric may require the ports connected by the
14 link to login for reassignment of temporary IDs to the ports. A fabric may also support a
15 “state change notification” process in which ports having operational links to the fabric
16 may request to be notified by the fabric when a state change is detected.

17
18 As described above, the data processing system in FIG. 1 has facilities for
19 ensuring that the logical storage volumes 28, 29, 30, and 31 are accessible to the hosts 22,
20 23, 24, 25. It is not only physically possible for any host to have access to any logical
21 volume in the data storage subsystem, but for fault tolerance is desirable for there to be at
22 least two independent paths through the data network 21 and through the cached storage

1 subsystem from each host to each logical storage volume.

2

3 Although it may be physically possible in the data processing system of FIG. 1 for
4 any host to have access to any logical storage volume, it may be desirable to restrict the
5 set of volumes that can be seen by any one host. Certain “private” volumes should be
6 assigned to each host and other hosts should not be permitted to see or modify the private
7 volumes of other hosts. Moreover, the “boot” process for a host is slowed down by
8 searching for and reporting all the volumes to which the host has access. Certain
9 operating systems are limited by the number of devices that they can handle at a given
10 period of time, and for a host using such an operating system, it is not only desirable but
11 also necessary to limit the number of volumes that the host can access.

12

13 The problem of restricting the set of volumes that can be seen by any one host is
14 solved by a method of named groups and/or a method of virtual ports. The method of
15 named groups is applicable to all Fiber Channel topologies for the connections of hosts to
16 one or more ports of a storage controller. The method of virtual ports causes at least one
17 port of the storage controller to appear as if it were a port of a fabric. The method of
18 virtual ports could cause the port of the storage controller to appear as if it were a Fibre-
19 Channel FL_Port (a port in a loop that is part of a fabric), E_Port (a port that is an
20 interlink between two switches), or F_Port (fabric port). The two methods could be used
21 at the same time in the same storage controller. In either the method of named groups or
22 the method of virtual ports, the limited set of volumes accessible to a host can be

1 specified by a list of logical volumes in the limited set, or by a procedure that defines the
2 limited set of volumes. An example of a specification for a procedure that defines the
3 limited set of volumes is a vector defining what will be called a “disk spread.”

4

5 Referring to FIG. 4, there are shown further details of the data processing system
6 of FIG. 1. In FIG. 4, each of the hosts 22, 23, 24 and 25 is shown to have a respective
7 host controller 61, 62, 63, 64 for each host port 65, 66, 67, 68 directly linked to the port
8 adapter 35. In the context of this specification, the term “host controller” refers to the
9 controller of the host ports on the data network. In the data processing art, such a host
10 controller is often referred to as a “host bus adapter” or “HBA”. The term “host
11 controller” will be used instead to avoid confusion with the “port adapter” of the storage
12 controller or storage subsystem. Each host controller 61, 62, 63, 64 is programmed with
13 a unique WWN for circuitry in the host controller for each of the respective ports 65, 66,
14 67, 68. In other words, if a host controller is replaced with a different host controller, the
15 WWN associated with the host port will change. Moreover, a host controller may
16 provide more than one port, and the host controller is programmed with a different WWN
17 for each of its ports.

18

19 As further shown in FIG. 4, the port adapter 35 includes port circuitry 71 and 72
20 for each of its two ports 73 and 74. The port circuitry, for example, includes application-
21 specific integrated circuits (ASICs) for communicating over the loops 41 and 43 in
22 accordance with the Fibre Channel standards. The port circuitry is connected to an

1 input/output bus 75 of a microprocessor 76. The microprocessor 76 is connected to a
2 random access memory 77 by a memory bus 78. The microprocessor 76 is also
3 interfaced to the storage controller busses 33 and 34. The memory 77 stores microcode
4 79 executed by the microprocessor 76.

5

6 STORAGE VOLUME PARTITIONING BY NAMED GROUPS

7

8 In order to restrict the set of volumes that can be seen by any one host, the
9 memory 77 of the port adapter 35 stores information defining a correspondence between
10 hosts in the data processing system and the set of volumes accessible to each host through
11 the port adapter. For example, a volume access table 80 and volume lists 81 are stored in
12 the memory 77. The volume access table specifies a correspondence between hosts and
13 respective lists of volumes accessible to the hosts. A back-up copy of the volume access
14 table and volume list could be stored in one of the storage volumes of the cached storage
15 subsystem.

16

17 The information in the volume access table 80 and the volume lists 81 can be
18 accessed by a system administrator viewing a display 91 and operating a keyboard 92 of a
19 service processor 93 interfaced to the controller busses 33, 34. The display 91, keyboard
20 92, and service processor 93, for example, are provided by a conventional lap-top
21 computer mounted behind a door (not shown) in the cabinet (not shown) of the storage
22 controller. The service processor 93, for example, is programmed to provide a graphical

1 user interface for volume configuration, partitioning, and other storage management
2 functions. The service processor also has a floppy-disc drive for down-loading of the
3 microcode 79 from a from a floppy-disc program carrier 95. Alternatively, a system
4 administrator at a remote terminal or host could access the information in the volume
5 access table 80 and volume lists 81 and down-load microcode via a modem or dedicated
6 link or via the data network (21 in FIG. 1) using an appropriate communications protocol
7 such as the Simple Network Management Protocol (SNMP).

8

9 Referring to FIG. 5, there is shown a simplified construction for a volume access
10 table 82. A preferred form of construction for the volume access table 80 in FIG. 4 will
11 be described below with reference to FIGs. 7 to 10. In the simplified construction of FIG.
12 5, each entry in the volume access table includes a volume group name, a host controller
13 port WWN, a host controller port S_ID, a private/shared flag, and a pointer to a volume
14 list. The volume group name includes a host name and a number identifying a particular
15 one of possibly many host ports linked to the port adapter. The volume group name
16 provides a stable and unique identification number for a group of logical storage volumes
17 to be accessed from a host port. At any given time, the S_ID and host controller port
18 WWN also provide unique identification numbers for the volume group, but the S_ID is
19 changed upon booting of the host controller, and the WWN is changed upon controller
20 replacement. Using a stable ID permits automation of booting, and simplifies
21 management of controller replacement. Moreover, by providing a path through the data
22 network from a host controller to each port adapter and by including an entry in the

1 volume access table in each port adapter defining a volume group name for the host
2 controller, it is possible to provide port independence in the sense that a host controller
3 may access its corresponding set of logical volumes through any of the port adapters.
4 When set, the private/shared flag indicates that all of the logical volumes in the volume
5 list are private and no permission is needed from a lock manager before the host
6 controller port can accessing any logical volume in its assigned volume list.

7
8 The volume group names and the volumes in each volume list are defined by the
9 system administrator. The host controller port WWNs are entered into the volume access
10 table 80 automatically during installation or replacement of a host controller. The host
11 controller S_ID is entered and updated automatically during each boot. The system
12 administrator defines the host names, for example, by sequentially assigning numbers to
13 the hosts. The host names are also known to the hosts, so that the relationship between
14 each host and the volumes assigned to the host can be re-established automatically if both
15 the S_IDs and WWNs would happen to change, for example as a result of host controller
16 replacement and a change in the configuration of the data network. In this situation, the
17 host controller port can transmit its corresponding group name to the storage subsystem
18 during a login process or in response to a request from the storage subsystem in response
19 to a state change notification so that the storage subsystem can reestablish the relationship
20 of the host controller port's volume group name and volume list with respect to its new
21 WWN and S_ID.

22

1 In the table 82 of FIG. 5, for example, the entry "HOST22-1" indicates the first
2 controller port of host no. 22, and it could be coded as a number 221, although the entry
3 could be displayed as "HOST22-1" for viewing by the system administrator. The other
4 entries in the table are shown as hexadecimal numbers. The volume access table itself
5 can have any suitable organization in the port adapter memory 77, such as a fixed-size
6 block of contiguous memory locations, or doubly-linked lists of the table entries.

7
8 The specification of the particular volumes in a volume group should permit
9 flexibility in assigning a variable number of volumes to each group, and also permit
10 flexibility in defining overlap between the groups corresponding to certain volumes that
11 are shared by the host processors. However, controller memory is a valuable resource,
12 and it is necessary to rapidly determine whether or not a volume specified by a storage
13 access request from a host controller port is included in the port's volume list. The
14 searching of a long list of pseudo-random logical volume numbers, for example, would
15 directly impact storage access time. Search time is especially critical in a cached storage
16 system, since the search time may be comparable to the data access time if the requested
17 data resides in the cache memory. In any case, the assignment of a long list of pseudo-
18 random logical volume numbers complicates the storage management problem, because it
19 is difficult for the storage administrator to comprehend how the volumes are assigned or
20 shared when the volumes in the volume list become nearly full and additional volumes
21 must be allocated to the hosts.

22

1 Hashing techniques can be used to search a long volume list rather efficiently, for
2 example, by using the last couple of digits of the specified logical volume number as a
3 hash table index. Since the list is changed infrequently, the list entries can be sorted for a
4 binary search procedure. Moreover, it is possible to reduce the number of entries in the
5 volume list and speed up the search process by specifying a range or vector of volume
6 numbers in each list entry. Shown in FIG. 6, for example, is a volume list entry 83
7 specifying a beginning volume number 84, an ending volume number 85, and a value 86
8 of one minus a stride (S). For example, the beginning volume number 84 is coded in
9 three bytes, the ending volume number is coded as three bytes, and the stride is coded as
10 two bytes. The stride (S) indicates the difference between neighboring volume numbers
11 of the volumes in a disk spread. The stride (S), for example, is a positive integer ranging
12 from 1 to 256 decimal.

13

14 Referring to FIG. 7, there is shown a preferred form of construction in which the
15 volume access table 80 in the memory 77 of the port adapter 35 stores only the volume
16 access information need for processing of a data access request by a host controller, and
17 an additional volume access table 88 in the cache memory 32 stores information that
18 appears in the table 82 of FIG. 5 but is not needed to process a data access request by a
19 host controller. The table 80 in the port adapter memory 77 is further identified as a
20 “transient” volume access table because it includes the transient S_IDs of the host
21 controller ports that have logged in to the port adapter 35. The table 88 in the cache
22 memory 32 is further identified as a “persistent” volume access table because it includes

1 the relatively permanent WWNs of the host controller ports that have logged in to the
2 port adapter 35. As further shown in FIG. 7, the port adapter 36 has a similar transient
3 volume access table 87 in its port adapter memory 90. The transient volume access table
4 87 includes the transient S_IDs of the host controller ports that have logged in to the port
5 adapter 36. In this case, example, the persistent volume access table 88 includes the
6 WWNs of the host controller ports that have logged in to the port adapter 90 in addition
7 to the WWNs of the host controller ports that have logged in to the port adapter 35. It is
8 desirable to share a persistent volume access table among a number of transient volume
9 access tables in the case where a host controller may access the same volume group from
10 the ports of different port adapters, since this avoids duplication of persistent volume
11 access table entries that would otherwise occur. Also shown in FIG. 7 are volume
12 attributes and locking information 89 stored in the cache memory 32, and respective table
13 directories 97, 98 that are stored in the port adapter memories 77, 90 and used to speed up
14 the process of searching the transient volume access tables 80 and 87.

15

16 The persistent volume access table 88 and the volume attributes and locking
17 information 89 also reside in a logical storage volume and be maintained in cache in a
18 fashion similar to any other logical storage volume accessed by a host controller.
19 However, it is preferred to locate the persistent volume access table 88 and the volume
20 attributes and locking information in a reserved area of cache memory, called "shared
21 memory," that is never deallocated, so that this information will always be found in the
22 cache memory when a port adapter attempts to access it.

1

2

3

4

5

6

7

8

9

10

11

12

13

14

15

16

17

18

19

20

21

22

Referring to FIG. 8, there is shown a preferred construction for the persistent volume access table 88. The persistent volume access table 88 includes an entry for each distinct group of volumes to be accessible to a particular host controller from a particular port adapter. The persistent volume access table 88 has columns for volume group names, host controller WWNs, a private/shared flag, and pointers to volume lists, and these columns are similar to the corresponding columns in the table 82 of FIG. 5. The persistent volume access table 88 further includes columns for a volume bitmap and a port adapter index list. A volume bitmap has a respective bit for each logical volume in the storage subsystem, and the respective bit is set if the logical volume associated with the bit is included in the volume list, and not set if the logical volume associated with the bit is excluded from the volume list. The port adapter index list includes the table index of each corresponding entry in a transient volume access table. Therefore, the port adapter index list could include an index for each port adapter to which the host controller is currently logged into. It may be desirable, but not necessary, for the hosts to log out of the port adapters. This would ensure that the volume access tables more precisely reflect the current state of the data network, and prevent any misdirection of messages if a network failure would prevent a storage adapter from being notified of a state change that would cause any S_ID in the transient volume access table to be reassigned.

In the persistent table 88 of FIG. 8, the information in each volume bitmap is redundant with the information in each volume list. In this case, the volume list is the

1 original form of the information as specified by the system administrator at configuration
2 time. The volume bitmap is used for rapidly searching whether a specified volume is in
3 the volume list, for example in response to a data access request from the host controller
4 port to which the volume list is associated. The volume bitmap is generated
5 automatically from the volume list when the volume list is entered or modified by the
6 system administrator. The volume list is retained for viewing and revision by the system
7 administrator. In a system with a limited number of logical volumes, the volume list
8 could be quickly generated from the volume bitmask, and therefore there would be no
9 need to retain the volume list in memory or include the volume list in the volume
10 access table.

11

12 Referring to FIG. 9, there is shown a preferred format for the transient volume
13 access table 80. The table 80 includes an entry for each host controller port currently
14 logged into the port adapter. The table 80 includes columns for the host controller S_ID,
15 the private/shared flag for the volume list, the volume bitmap, and a persistent table
16 index. The persistent table index in each entry of the transient volume access table 80
17 points to the entry in the persistent volume access table having the same private/shared
18 flag and volume bitmap as the entry of the transient volume access table, and having a
19 host controller WWN corresponding to the host controller S_ID in the entry of the
20 transient volume access table.

21

22 Referring to FIG. 10, there is shown a preferred format for the table directory 93.

1 The table directory includes, in each entry, a list of pairs of the S_ID and corresponding
2 index to the transient volume access table for each SID having a certain modulus of the
3 S_ID. The modulus each S_ID in the table directory corresponds to the index of the
4 directory table entry containing the S_ID. For example, the modulus of the S_ID is a
5 certain number of least significant bits of the S_ID, and in this case the number of entries
6 in the table directory 93 is two raised to the number of the least significant bits.

7
8 To search the transient volume access table 80 for an entry having a specified host
9 controller S_ID, an index is computed by masking off the least significant bits of the
10 specified S_ID to obtain the corresponding modulus, and the modulus is used to index the
11 table directory to find a list of S_IDs and corresponding indices that should include the
12 specified S_ID if the specified S_ID is somewhere in the transient volume access table.
13 The list in the entry indexed by the modulus is searched for the specified SID, and if it is
14 found, then its corresponding index is the index to the transient volume access table entry
15 containing the specified S_ID. If the specified SID is not found in the list in the entry
16 indexed by the modulus, the specified S_ID does not appear in the S_ID column of the
17 transient volume access table.

18
19 Since S_IDs are assigned to hosts controller ports sequentially when the host
20 controller ports share a loop, the size of the table directory 93 can be determined based on
21 reasonable constraints on the network geometry. For example, if the most complex
22 network geometry connects a certain maximum number of loops to one or both of the two

1 physical ports of the port adapter (for example, as described below with reference to FIG.
2 38), then the number of entries in the table directory can be chosen to be greater or equal
3 to the maximum number of host controller ports on a single loop, and the number of pairs
4 of S_IDs and corresponding transient table indices in each entry of the table directory can
5 be set to the maximum number of loops that can be connected to each port through one
6 switch. However, the number of possible S_IDs in each entry of the table directory 97
7 can be reduced still further if one assumes a particular kind of connection between each
8 loop and each port adapter port. For example, if a single switch is used to connect loops
9 to the port adapter (for example, as shown in FIG. 38), and all of the host controller ports
10 are assigned S_IDs sequentially by this switch, then the number of entries in the table
11 directory can be chosen to be the smallest power of two greater or equal to the maximum
12 number of host controller ports linked to the port adapter ports, and there will be at most
13 one S_ID and corresponding transient volume access table index in each entry of the table
14 directory.

15

16 Referring to FIG. 11, there is shown a flowchart of a microcode routine executed
17 by the microprocessor (76 in FIG. 4) of the port adapter (35 in FIG. 4) in response to a
18 "Report LUNs" or access request by a host. In a first step 101, the microprocessor
19 decodes the S_ID from the request. Then in step 102, the microprocessor searches the
20 volume access table for an entry having the S_ID from the request. For example, this is
21 done by indexing the table directory of FIG. 10 with a modulus of the S_ID from the
22 request, searching the indexed entry of the table directory for the S_ID from the request,

1 and if the S_ID from the request is found in the entry, then indexing the volume access
2 table of FIG. 9 with the corresponding index in the entry of the table directory. If the
3 S_ID of the request is not found in the volume access table, then execution branches to
4 step 104 to report no LUNs or deny the access request, and the routine is finished.

5

6 If the S_ID of the request is found in the volume access table, then execution
7 continues to step 105. In step 105, the volume bitmap or volume list is accessed. For the
8 volume access table format of FIG. 9 the volume bitmap is accessed, and for the volume
9 access table format of FIG. 5, the volume list is accessed. In step 106, execution
10 branches to step 107 for a "Report LUNs" request. In step 107, the port adapter reports to
11 the host controller a storage LUN for each volume in the volume bitmap or volume list,
12 and the routine is finished. For an access request, execution continues from step 106 to
13 step 108. In step 108, execution branches to step 109 if the volume to access is not in the
14 volume bitmap or volume list. For a bitmap, this can be done simply by addressing the
15 bit corresponding to the logical volume specified by the host controller. In general, the
16 volume list is searched by using a technique most suitable for quickly searching the
17 volume list. In step 109, the port adapter denies the storage controller access to the
18 logical volume, and the routine is finished.

19

20 If the volume to access is in the volume bitmap or list, then execution continues
21 from step 108 to step 110. In step 110, the private/shared flag is inspected for the
22 indexed entry of the volume access table. If the private/shared flag is set, then execution

1 continues from step 110 to step 111. In step 111, the port adapter accesses the logical
2 volume specified by the host controller, and the routine is finished. If the private/shared
3 flag is not set, then execution branches from step 110 to step 112 to access the locking
4 information the cache memory. If the private/shared flag is set, then access is permitted,
5 and execution branches from step 113 to step 111 to access the volume. If the volume to
6 access is public and is already locked in a fashion incompatible with the access requested
7 by the host (e.g., the volume is already write locked and the host controller requests a
8 read or a write access, or the volume is already read locked and the host controller
9 requests a write access) then access is not presently permitted. The host controller's
10 S_ID is placed on a wait list, in order to notify the host controller when the logical
11 volume becomes available; execution branches from step 113 to step 114 to temporarily
12 deny access to the volume, and the routine is finished. If the volume to access is public
13 and is not locked or is locked in a fashion compatible with the requested access by the
14 host controller (e.g., the volume is already read locked and the host controller requests a
15 read access), then a lock is granted to the host controller, and execution branches from
16 step 113 to step 111 to access the volume, and the routine is finished.

17

18 Referring to FIG. 12, there is shown a flowchart 120 of a microcode routine 120
19 executed by the microprocessor (76 in FIG. 4) of the port adapter (35 in FIG. 4) when
20 converting the vector representation of FIG. 6 to a list of logical storage volume
21 identifiers. This routine is used, for example, in step 106 to report the storage LUNs in a
22 disk spread defined by an entry in the volume list. In a first step 121, an index (I) is set to

1 the first volume number (BEGIN) for the disk spread. Next, in step 122, the value of the
2 index I is inserted onto the list. Then in step 123, the value of the index is compared to
3 the last volume number (END) of the disk spread. If the value of the index is greater or
4 equal to the value of the last volume number (END), then the routine terminates.
5 Otherwise, execution continues from step 123 to step 124. In step 124, the index (I) is
6 incremented by the stride (S), and execution loops back to step 122.

7
8 Referring to FIG. 13, there is shown a flowchart 130 of a microcode routine
9 executed by the microprocessor (76 in FIG. 4) of the port adapter (35 in FIG. 4) when
10 determining whether or not a specified logical storage volume is included in a disk spread
11 defined by the vector representation of FIG. 6. This routine is used, for example, in step
12 107 to determine whether or not the volume specified by a data access request is included
13 in a disk spread defined by an entry of the volume list. In a first step 131, the volume
14 index (I) is compared to the first logical volume number (BEGIN) of the disk spread.

15
16 If the volume index (I) is less than the first logical volume number, the execution
17 returns indicating that the specified logical storage volume is not included in the disk
18 spread. Otherwise, execution continues from step 131 to step 132. In step 132, the
19 volume index (I) is compared to the last logical volume number (END) of the disk spread.
20 If the volume index (I) is greater than the last logical volume number, then execution
21 returns indicating that the specified logical storage volume is not included in the disk
22 spread. Otherwise, execution continues from step 132 to step 133.

1

2 In step 133, the stride (S) is compared to one. If the stride (S) is equal to one,
3 then execution returns indicating that the specified logical storage volume is not included
4 in the disk spread. Otherwise, execution continues to step 134. In step 134, the
5 difference is calculated between the index I and the first logical volume number (BEGIN)
6 of the disk spread. Then a value REM is computed which is the modulus of the
7 difference (DIF) and the stride (S). In particular, if the difference (DIF) is zero, then
8 REM is zero. Otherwise, DIF is divided by S using an integer division procedure, and
9 REM is the remainder of the division. If REM is zero, then execution returns indicating
10 that the specified logical volume is in the disk spread; otherwise, execution returns
11 indicating that the specified logical volume is not in the disk spread.

12

13 Referring to FIG. 14, there is shown a graphical display 140 of disk spreads in a
14 two-dimensional volume space. For example, the logical volumes I0 to I12 are arranged
15 as a four-by-four square matrix. The rows of the matrix correspond to respective disk
16 spreads defined by the respective vectors $V0 = (0, 3, 1)$, $V1 = (4, 7, 1)$, $V2 = (8, 11, 1)$,
17 and $V3 = (12, 15, 1)$. For each of these row vectors V0 to V3, the stride (S) is one. Also
18 shown is a disk spread defined by a column vector $V4 = (2, 14, 4)$ having a stride (S) of
19 four. In a multi-processor system, for example, the disk spreads represented by each of
20 the row vectors V0, V1, V2, V3 could be assigned to a respective one of four host
21 processors, and the disk spread represented by the column vector V4 could be assigned to
22 a fifth host processor. In this case, none of the first four processors would share any of

1 the volumes with any other of the first four processors, but each of the first four
2 processors would share one volume with the fifth processor. In this fashion, a limited
3 amount of information can be stored in the memory of a port adapter to define a rather
4 large number of volumes having various shared and private relationships among a
5 multiplicity of host processors.

6

7 Referring to FIG. 15, there is shown a graphical display of disk spreads in a three-
8 dimensional volume space. The volume space includes 4,096 volumes, or 16 volumes
9 along each of the x, y, and z axes. Shown in this volume space is a first disk spread
10 defined by a vector $V5 = (0, 15, 1)$ having a stride (S) of one, a second disk spread
11 defined by a vector $(15, 255, 16)$ having a stride (S) of sixteen, and a third disk spread
12 defined by a vector $V7 = (255, 4095, 256)$. By using such a three-dimensional graphical
13 display, it is possible to show complicated private and shared relationships between the
14 volumes accessible to a large number of host computers. For example, the service
15 processor 93 in FIG 4 could provide such a three-dimensional view on the display. The
16 volumes associated with a particular group name could be displayed in a corresponding
17 primary color, so that the private and sharing relationships for up to three selected group
18 names or host controller ports could be displayed simultaneously. For example, in FIG.
19 15, the disk spread specified by the vector V5 could be associated with a first host
20 processor and displayed in red, the disk spread specified by the vector V6 could be
21 associated with a second host processor and displayed in blue, and the disk spread
22 specified by the vector V7 could be associated with a third host processor and displayed

1 in green. Therefore the shared volume I15 would be displayed in magenta, and the shared
2 volume I255 would be displayed in cyan. Moreover, the service processor could be
3 programmed to display a colored two-dimensional 16 x 16 square matrix of the logical
4 volumes for any selected one of 16 planes parallel to the x, y, or z axes in the logical
5 volume space. In addition to indicating the private and shared relationships of the
6 volumes to the hosts, the two-dimensional displays could provide additional information
7 in the display region of each volume, such as a numeral or graph indicating of the free
8 storage space in each volume.

9

10 Referring to FIG. 16, there is shown a flowchart 160 of a microcode routine
11 executed by the microprocessor (76 in FIG. 4) of the port adapter (35 in FIG. 4) for
12 volume configuration when installing the storage subsystem into the data processing
13 system of FIG. 1. In a first step 160, the system administrator enters a group name for a
14 host to be permitted access through one or more port adapters. For example, the system
15 administrator manually operates the keyboard 92 in Fig. 4 to enter the host name, host
16 controller number(s), and port adapter number(s). The service processor 93 forwards this
17 information to the port adapter(s) and the microprocessor(s) in the port adapter(s)
18 allocate(s) a table entry for each host controller number and places the respective group
19 name into each table entry. Then in step 162, the port adapter(s) automatically search the
20 network for the S_ID and WWN of the host controller ports corresponding to the group
21 names. For example, the port adapters poll the ports of the hosts, and the host controllers
22 are programmed to respond to port adapters by returning their host names and controller

1 numbers. If the polling is successful in returning the S_ID and WWN associated with the
2 group name, as tested in step 163, then execution continues to step 164. In step 164 the
3 system administrator enters a volume list for the group name, for example by entering
4 logical volume numbers, ranges of logical volume numbers, or vector specifications
5 (beginning number, ending number, stride) at the keyboard 92 in FIG. 4. In step 165, the
6 microprocessor creates a table entry for the group name and its associated S_ID, WWN,
7 and a pointer to the volume list. In step 166, execution continues to step 167 if the
8 system administrator has no more group names to enter; otherwise, execution loop back
9 to step 161. In step 167, the volume access table and volume lists are copied to a storage
10 volume as back-up for port adapter error recovery or diagnostic purposes, and the routine
11 if finished. The back-up copy should be updated if additional group names are added to
12 the volume access table or if the volume lists are changed

13

14 If the polling is not successful in returning a S_ID and WWN for a host or group
15 name, then execution branches from step 163 to step 168. In step 168, the items not
16 found are reported to the system administrator, who has the option of creating or not
17 creating a table entry if the items cannot be found. For example, it is possible for the
18 system administrator to permit a table entry to be created for host controllers that are not
19 presently logged into the data network, and the port adapter can fill in the missing items
20 automatically if and when the host controller are booted. If the system administrator has
21 selected the option of not creating a table when items are missing, then execution
22 branches from step 169 to step 166. Otherwise, execution continues from step 169 to step

1 170. In step 170, the S_ID and WWN for the group name is set to null, and then
2 execution continues to step 164.

3

4 Referring to FIG. 17, there is shown a flowchart 180 of a microcode routine
5 executed by the microprocessor (76 in FIG. 4) of the port adapter (35 in FIG. 4) when
6 notified of a network state change such as a host controller boot. In a first step 181, the
7 port adapter obtains the WWN, the S_ID, and the group name of the host controller port.
8 Then in step 182, the volume access table is indexed with the group name. If no entry for
9 the group name is found in the table, then the routine is finished. Otherwise, execution
10 continues to step 183. In step 183, the WWN obtained from the host controller port is
11 compared to the WWN in the table entry. If the WWN obtained from the host controller
12 port is the same as the WWN in the table entry, then execution continues from step 183 to
13 step 184. In step 184, the S_ID value in the table entry is reset to the value obtained by
14 the port adapter in step 181, and the routine is finished.

15

16 If in step 183 the WWN obtained from the host controller port is different from
17 the WWN in the table entry, then execution branches from step 183 to step 186. In step
18 186, execution branches to step 188 if the WWN in the table entry is null. The null value
19 in the table indicates that the system administrator has authorized the WWN to be set
20 during a controller boot, for example just after an authorized installation or change of a
21 host controller circuit board. Therefore, in step 188 the WWN and S_ID obtained in step
22 181 are entered into the entry of the volume access table, and the routine is finished. In

1 step 186, if the WWN entry is not null, then an unauthorized change in a host controller
2 circuit board has occurred, and execution branches from step 186 to step 187 to report the
3 error to the system administrator, and then the routine is finished.

4
5 Referring to FIG. 18, there is shown is a flowchart of the procedure that should be
6 performed by a system administrator during host controller replacement. In step 201 the
7 host is interrupted to suspend any communication over the data network or shut down to
8 permit replacement of the host controller. Unless the host controller board has a “hot
9 swap” capability, the host will be shut down to shut off power to the controller board
10 before it is removed and replaced. Then in step 202, the system administrator nullifies
11 the host controller’s WWN in all such WWN entries in all volume access tables in all port
12 adapters of the storage subsystem, and in step 203 the host controller is replaced. The
13 replacement of the host controller board causes the WWNs of the ports of the host
14 controller board to change. In step 204, host processing is restarted or continued,
15 including a boot of the new host controller board, and the procedure is finished.

16

17 MAPPING OF LUNs TO LOGICAL VOLUME NUMBERS

18

19 As described above with reference to FIG. 7, a host may request access to a
20 specified logical volume in the storage subsystem. If a host may access only a very
21 limited set of the logical volumes in the data storage subsystem, however, it may be
22 desirable for the port adapter to report back to each host a limited range of LUNs

1 representing logical volumes that the host can access, and for the port adapter to map this
2 limited range of LUNs to the limited set of logical volumes that each host can access.
3 The host therefore does not need to deal with a large number of possible volumes, and
4 minimal modification of the host programming is needed when a host is networked to the
5 data storage subsystem. Instead of specifying a logical volume number when requesting
6 access to a logical storage volume, the host specifies a LUN, and the port adapter
7 receiving the storage access request translates the specified LUN to a corresponding
8 logical volume number. There can be a different mapping of LUNs to logical volume
9 numbers for each host or host controller port.

10

11 As shown in FIG. 19, the mapping of LUNs to logical volume numbers for a
12 volume group name is specified by the entries in the volume list 210 associated with the
13 volume group name. Each entry in the volume list 210 includes a LUN in the volume
14 group and its corresponding logical volume number (VOL_NO.). The first entry 211 in
15 the volume list, for example, maps LUN 1 to logical volume number 8.

16

17 In practice, it is desirable to assign ranges of contiguous LUNs to the volume
18 groups or hosts. These ranges of contiguous LUNs, for example, are evident when the
19 LUNs in the volume list are sorted by LUN, for example as shown in FIG. 19. In many
20 cases, each range of contiguous LUNs corresponds to a range or vector of logical volume
21 numbers. Therefore, in a data processing system having a large number of LUNs per host
22 and an even larger number of logical volumes, it may be desirable to specify the LUN to

1 logical mapping as a number of ranges of contiguous LUNs and corresponding vectors of
2 logical storage volume numbers.

3

4 As shown in FIG. 20, for example, each entry in the volume list 220 associates a
5 respective range of contiguous LUNs with a vector of logical volume numbers. Each
6 entry includes the first LUN in the range, the number (N) of LUNs in the range, the
7 beginning logical volume number in the associated vector of logical volume numbers,
8 and the stride (S) of the vector. The first entry 221 in the volume list 220, for example,
9 specifies a first LUN of one, a number (N) of three, a beginning logical volume number
10 of eight, and a stride (S) of two. This maps the sequence of contiguous LUNs (1, 2, 3) to
11 the sequence of logical volume numbers (8, 10, 12), and therefore the first entry 221 in
12 the volume list 220 of FIG. 6 encodes the same information as the first three entries in the
13 volume list 210 of FIG. 19.

14

15 To find whether a specified LUN is within the map of any entry in the volume list
16 220, the specified LUN value is compared to the LUN value in the entry, and if it is less
17 than the LUN value in the entry, the specified LUN is not within the map of the entry. If
18 the specified LUN value is greater or equal to the LUN value of the entry, then the
19 specified LUN value is compared to the LUN value in the entry plus the value of N in the
20 entry. If the specified LUN value is less than the LUN value in the entry plus the value of
21 N in the entry, then the specified LUN is within the map of the entry; otherwise, it is not.
22 If the specified LUN is within the map of the entry, then its corresponding logical storage

1 volume number is computed as:

2

3
$$\text{LOGICAL_VOLUME_NO.} = \text{MAP_VOL_NO.} + (\text{SPEC_LUN} - \text{MAP_LUN}) * S$$

4

5 where MAP_VOL_NO. is the VOL_NO in the map entry, SPEC_LUN is the
6 specified LUN value, and MAP_LUN is the LUN value in the map entry.

7

8

9 VOLUME PARTITIONING BY VIRTUAL PORTS

10

11 As described above with reference to FIG. 7, the port adapters can be programmed
12 to provide a “Report LUNs” command that will report to each host only the LUNs or
13 logical storage volumes assigned to the host. In other words, a host would not normally
14 have the capability of seeing the LUNs or logical volumes assigned to another host. This
15 offers some security, but it is incompatible with the conventional “Report LUNs”
16 command of the Fibre Channel standards. In a Fibre Channel network, if a conventional
17 “Report LUNs” command is used to discover LUNs, all ports will report back all of their
18 LUNs.

19

20 The method of FIG. 7 also has the disadvantage that a significant amount of port
21 adapter processing time is spent in managing variable-length volume lists or very long
22 bitmaps. The method of FIG. 7 further has the disadvantage that the entire group of

1 volumes for a host controller must be flagged as shared if only one of the volumes in the
2 group are shared, which increases the access time of the private volumes in the group. It
3 would be desirable to use a predefined, fixed-length format for the lists that are frequently
4 accessed if this would not waste port adapter memory or unduly limit the set of logical
5 volumes that could be assigned to each host. For a host controller having both shared and
6 private volumes, it would be desirable to associate more than one group of volumes to the
7 host controller, including one or more groups of all private volumes and one or more
8 groups of all public volumes, so that each group can be separately flagged as private or
9 public.

10

11 The method of virtual ports overcomes these disadvantages in a way that is
12 compatible with the Fibre Channel specifications. In accordance with the method of
13 virtual ports, the storage subsystem presents to the Fibre Channel network a set of
14 “virtual” Fibre Channel ports that do not really exist on the network. A set of logical
15 volumes is assigned to each of the virtual ports. The logical volumes within each set are
16 accessible from the virtual port through at least one physical port of the storage
17 subsystem. This physical port is therefore a fabric port and the storage subsystem
18 provides a virtual switch from the physical port to each of the virtual ports accessible
19 through the physical port. In particular, the physical port is a Fibre Channel FL_Port if it
20 is in a loop of the Fibre Channel network, or it is a Fibre Channel E_Port if it is
21 connected to a switch, or otherwise it is a Fibre Channel F_Port. The port adapter
22 providing the physical port is programmed to function as an FL_Port, E_Port or F_Port

1 by responding to host login commands and assigning S_IDs; by responding to
2 conventional Fibre Channel Report LUNs commands for reporting LUNs assigned to the
3 virtual ports; and by responding to routing instructions from a host for routing data access
4 requests to a specified virtual port and a specified LUN of the virtual port.

5
6 When using the method of virtual ports for volume partitioning, one or more
7 virtual ports are assigned by the system administrator to each of the hosts having access
8 to one or more of the virtual ports. Preferably logical storage volumes can be accessed
9 through a single virtual port by no more than one assigned host. This simplifies the
10 access control or filtering function of the storage subsystem, and does not unduly restrict
11 the sets of logical volumes that can be accessed since a sufficiently large number of
12 virtual ports can be created. A fixed-length format can be used for storing the lists of
13 hosts, virtual ports, and LUN to logical storage volume mapping information for each
14 virtual port because a variable number of virtual ports can be assigned to each host. If a
15 host needs more logical storage volumes than can be included in the volume list for a
16 single virtual port, another virtual port can be created and assigned to the host. Moreover,
17 each virtual port can be accessible through any desired number of the storage subsystem
18 adapter ports defined as FL_Ports, E_Ports, or F_Ports. The virtual port is made
19 accessible by storing in the memory of the port adapter the access and mapping
20 information for the virtual port.

21
22 A host uses conventional Fibre Channel commands and protocols for sending

1 requests to its assigned virtual port or ports. Since the Fibre Channel standards do not
2 restrict the ability of a host to send requests to all of the ports linked by the Fibre Channel
3 network, there must be a mechanism for the system administrator's assignment of the
4 virtual ports to the hosts to be communicated to the hosts, and the hosts must use this
5 information to direct their storage access requests to the virtual ports to which they are
6 assigned. This assignment information must also be used by the host if the host has an
7 operating system that permits the host to boot from a logical volume in storage linked by
8 the Fibre Channel network to the host, or that permits the operating system of the host to
9 collect information about the logical storage volumes that it can access. In other words,
10 the operating system routine that searches for the storage volumes that are accessible to
11 the host must send Report LUNs commands to only the virtual ports assigned to the host
12 and not to the virtual ports assigned to other hosts.

13

14 With reference to FIG. 21, the port adapter 36 of the cached storage subsystem 20
15 is shown having been programmed to provide virtual ports for access to respective groups
16 of logical storage volumes. The port adapter 36 has two physical ports 231 and 232
17 provided by port circuitry 233 and 234, respectively. The port circuitry performs
18 serializing, framing, sequencing, flow control, and coordination of protocols and services
19 in accordance with the Fibre Channel standards. The port adapter 36 further includes a
20 microprocessor 235 and a random access memory 236. The microprocessor 235 executes
21 microcode 237 in the random access memory 236. In particular, the microprocessor 235
22 is instructed by the microcode 236 to handle Fibre Channel requests received by the port

1 circuitry 233 in such a way that there appears to be a respective switch 238, 239 in the
2 storage controller 27 connecting each of the ports 123, 131 to a respective set of ports
3 240, 241 between the ports 123, 131 and the storage volumes (28, 29, 30, and 31) in the
4 cached storage subsystem 20. Since the ports 240, 241 do not appear as physical nodes
5 anywhere in the cached storage subsystem 20, the ports 240, 341 are called virtual ports,
6 and the switches 238 and 239 are called virtual switches.

7

8 As far as the data network 21 is concerned, the virtual ports 240 and 241 function
9 in a fashion similar to physical ports in the data network. The virtual ports are assigned
10 respective port names, WWNs, and S_IDs, and requests from other nodes in the network
11 are routed to them through their respective virtual switches 238, 239. For example, the
12 virtual ports respond to Report LUN requests in a fashion similar to physical ports. In
13 order to define the structure of the virtual switches, the random access memory 236 is
14 programmed with switch definitions 242 and also stores switch state information 243.
15 The switch definitions 242, for example, include the respective WWNs of the virtual
16 switches 238 and 239, and the switch state 243 includes the S_IDs assigned at any given
17 time to the virtual ports 240 and 241. The microcode 237 is also programmed to provide
18 respective name servers 244 and 245 which provide nodes linked to the respective ports
19 231 and 232 with limited access to the switch definitions 242 and switch state 243. In
20 particular, the name servers can provide the respective names, WWNs, and S_IDs of all
21 the virtual and physical ports of the virtual switches.

22

1 The random access memory 236 is further programmed with volume access and
2 mapping information 246. The volume access information indicates a respective set of
3 storage LUNs that any host can access by requests addressed to the virtual ports 240, 241.
4 The volume mapping information indicates respective logical volumes mapped to the
5 virtual ports 240, 241. The volume access and mapping information for each of the
6 virtual switches 238, 239 could be the same or it could be different. For example, it
7 could be the same to permit any host to be disconnected from any one of the network
8 loops 42, 44 and connected to the other loop without any change in the storage access
9 privileges or procedures of the host. The set of logical storage volumes mapped to
10 different virtual ports can be the same or different. For example, the set of logical storage
11 volumes mapped to two virtual ports can be the same to permit two hosts to share the set
12 of logical storage volumes mapped to the two virtual ports.

13
14 With reference to FIG. 22, there is shown a data processing system having a
15 cached storage subsystem 250 and four host computers 251, 252, 253, 254 linked to the
16 cached storage subsystem by a data network 255. In this example, the hosts are linked by
17 four loops 256, 257, 258, 259 to respective network port adapters 260 and 261 of the
18 cached data storage subsystem. Each port adapter has two physical ports designated as
19 “A” or “B”. The physical port 262 of the port adapter 260 is an “A” port linked to the
20 loop 256, the physical port 263 of port adapter 260 is a B port linked to the loop 257, the
21 physical port 264 of port adapter 261 is an “A” port linked to the loop 258, and the
22 physical port 265 of port adapter 261 is a “B” port linked to the loop 259. The port

1 adapters 260, 261 are programmed to provide respective virtual switches 266, 267 linking
2 their "A" and "B" physical ports to a set of virtual ports 268, including a respective
3 virtual port VP1, VP2, VP3, VP4 assigned to each of the hosts 251, 252, 253, 254. Each
4 of the port adapters 260 and 261 is programmed with respective volume access
5 information 269, 270 which can be identical to permit any of the hosts 251, 252, 253, 254
6 to be disconnected from any one of the loops 156, 157, 158, 158 and reconnected to any
7 other of the loops without any change in the storage access privileges or procedures of the
8 host.

9
10 With reference to FIG. 23, there is shown an example of the volume access and
11 mapping information 269 of FIG. 18. The volume access and mapping information 246
12 of FIG. 21 can have a similar format. The volume access and mapping information 269
13 includes a virtual port host table 281 listing each host having access rights through a
14 virtual switch controlled with the volume access and mapping information, and a virtual
15 port mapping table 282 listing each virtual port accessible through the virtual switch
16 controlled with the volume access and mapping information. Separate tables are used
17 because each host listed in the host table can have more than one assigned virtual port.
18 Also included in the volume access and mapping information 269 are optional lists 283 of
19 indices to the virtual port identifiers in the virtual port mapping table 282 assigned to
20 each host in the virtual port host table 281. The lists 283 are optional because the
21 information in the lists could be obtained by scanning through an entire column of host
22 indices in the virtual port mapping table. The lists 283 are not needed for handling a

1 storage access request by a host. The lists 283 are used for quickly responding to a host
2 when the host requests a report of the virtual ports from which it can access logical
3 storage volumes. It is desirable to compile the lists by scanning the host index column of
4 the virtual port mapping table and store the lists if a large number of hosts will be
5 requesting the information frequently, for example, if the hosts login to the network
6 frequently and request the information during each network login. The lists 283 in total
7 have a limited memory requirement because each virtual port in the virtual port mapping
8 table can have only one assigned host, so that in total the lists 283 include no more than
9 one occurrence of a virtual port index for each of the virtual ports listed in the virtual port
10 mapping table.

11

12 Referring to FIG. 24, there is shown a format for the virtual port host table 281.
13 The table 282 includes a unique row for each host controller port from which a host can
14 access a virtual port defined in the virtual port mapping table (282 of FIG. 23). The table
15 281 includes a column for the HOST_ID of the host controller port, a flag PORT A/B
16 indicating whether the host controller port is linked to the physical "A" or "B" port of the
17 port adapter programmed with the table 281, the WWN of the host controller port, the
18 S_ID of the host controller port, a V_PORT INDEX which is a first index to a row of
19 information in the virtual port mapping table (282 in FIG. 23) to which the host controller
20 port has access through the port adapter, and a V_PORT LIST POINTER which is zero if
21 there is only one V_PORT associated with the host controller port and otherwise points to
22 a list of additional virtual port mapping table indices to rows of information about

1 additional virtual ports to which the host controller has access through the port adapter.

2

3 The HOST_ID, for example, is the IP address of the host controller port. The
4 HOST_ID could have the same format as the volume group name of the table in FIG. 5;
5 i.e., a host name and controller sequence number

6

7 If the optional lists of indices to the virtual port table are not used, the V_PORT
8 LIST POINTER would simply be a flag indicating whether or not there is more than one
9 virtual port associated with the host controller port. If there is only one, it is found in the
10 V_PORT INDEX column of the virtual port host table. If there is more than one, then the
11 entire list can be found by scanning the HOST INDEX column of the virtual port
12 mapping table 282 for entries having a HOST INDEX matching the index of the row in
13 the virtual port host table containing the information for the host controller port.

14

15 Referring to FIG. 25, there is shown a format for the virtual port mapping table
16 282. The table 282 includes a unique row for each virtual port accessed through the port
17 adapter. The table 282 includes a column for the V_PORT_ID of the virtual port, a
18 HOST INDEX to a row of the virtual port host table for the host controller port associated
19 with the virtual port, a private/shared flag, and a LUN TO LOGICAL VOLUME MAP
20 specifying the set of LUNs at the virtual port, the set of logical storage volumes
21 accessible from the virtual port, and the mapping between the LUNs and the logical
22 storage volumes accessible from the virtual port. The V_PORT_ID is a primary key for

1 the table 282. The table 282, for example, is sorted by the V_PORT_ID or has an
2 associated hash table directory to facilitate searching of the table for a specified value of
3 the V_PORT_ID. The VPORT_ID could be an IP address for the virtual port.

4
5 The LUN TO LOGICAL VOLUME MAP in each row of the table 282 could use
6 the format of FIG. 19 or FIG. 20 with a fixed number of either the entries shown in FIG.
7 19 or the entries shown in FIG. 20. Since there would be a fixed number of possible
8 entries of the format of FIG. 19 or FIG. 20 but less than all of the entries might be used at
9 any given time, the unused entries could be filled in with a null value for the LUN, such
10 as a zero LUN value. The first entry would always be used, and any other entries used
11 would be contiguous with the first entry, so that all of the entries used could be found by
12 scanning and testing the LUN value in the entry. Alternatively, the LUN TO LOGICAL
13 VOLUME MAP for each virtual port could be preceded with an integer value indicating
14 the number of entries used in the map at any given time.

15
16 The LUN TO LOGICAL VOLUME MAP in each row of the table 282 could be a
17 bitmap in the form of a LUN to logical volume mapping matrix. The logical OR of all
18 the rows of such a matrix would be a bitmap indicating all of the logical volumes
19 accessible from the virtual port. The logical OR of all the columns of such a matrix
20 would be a bitmap indicating all of the LUNs accessible from the virtual port. The use of
21 such a mapping matrix would not be a practical alternative if there were a large number
22 of virtual ports, hosts, possible LUNs, and logical volumes, since the number of bits of

1 memory for storing such all of the mapping matrices would be the product o the number
2 of virtual hosts, the number of possible LUNs, and the number of logical volumes.

3

4 Referring to FIG. 26, there is shown a flowchart of a routine executed by the
5 microprocessor of the port adapter to provide a service for reporting the virtual ports
6 accessible to a host controller port through the port adapter. This service could provide
7 such a report when a host controller logs in to the data network, or when a host controller
8 port requests the information from a name server of the virtual switch provided by the
9 port adapter. In a first step 301, the microprocessor searches the virtual port host table for
10 an entry including a specified HOST_ID. In step 302, execution branches to step 303 if a
11 table entry is not found for the specified HOST_ID. In step 303, the port adapter reports
12 that there are no V_PORTS accessible to the host through the port adapter, and the
13 routine is finished.

14

15 In step 302, execution continues to step 304 if a table entry is found for the
16 specified HOST_ID. In step 304, the port adapter reports that the V_PORT of the
17 V_PORT index in the table entry is accessible to the host. In other words, the
18 microprocessor reads the V_PORT INDEX from the row of the virtual port host table
19 containing the specified HOST_ID, and then indexes the virtual port mapping table with
20 the V_PORT INDEX to find the first V_PORT_ID associated with the HOST_ID. In
21 step 305, the V_PORT LIST POINTER is compared to zero, and if it is zero, then the
22 routine is finished. Otherwise, execution continues from step 305 to step 306.

1

2 In step 306, the V_PORT list pointer is used to access the list of additional indices
3 to accessible V_PORTS in the virtual port mapping table. Finally, in step 307, the port
4 adapter reports that the indexed V_PORTS are accessible to the host through the port
5 adapter, and the routine is finished.

6

7 Referring to FIG. 27, there is shown a flowchart 320 of a routine executed by the
8 microprocessor of a port adapter when the port adapter receives from a host controller
9 port a volume access or Report LUNs request to a specified virtual port. To ensure
10 compliance with the Fibre Channel standards, this Report LUNs request might be
11 different from and in addition to the standard Report LUNs command. The Report LUNs
12 request of FIG. 27 reports the storage LUNs accessible to a specified host controller port,
13 in contrast to the standard Report LUNs command which could report all LUNs
14 accessible from a virtual port including all storage LUNs accessible from the virtual port
15 by any of the hosts. In response to a standard Report LUNs command directed to a
16 virtual port, the port adapter would search the virtual port mapping table 282 of FIG. 25
17 for the V_PORT_ID specified by the Report LUNs command to find the table entry
18 including the V_PORT_ID, and then report all of the LUNs included in the LUN TO
19 LOGICAL VOLUME MAP in the table entry.

20

21 In a first step 321 of the flowchart 320, the microprocessor searches the virtual
22 port mapping table for the entry of the virtual port specified by the host. If such a table

1 entry is not found, then execution branches from step 322 to step 323. In step 323, the
2 request is denied, and the routine is finished. If such a table entry is found, then
3 execution continues from step 322 to step 324. In step 324, the microprocessor obtains
4 the HOST_INDEX of the entry of the virtual port mapping table. In step 325, the virtual
5 port host table is indexed with the HOST_INDEX.

6
7 In step 326, the HOST_ID of the host controller port requesting access is
8 compared to the HOST_ID in the indexed entry of the virtual port host table. If the
9 HOST_ID of the host controller port requesting access is not equal to the HOST_ID in
10 the indexed entry of the virtual port host table, then execution branches from step 326 to
11 step 323. In step 323 the port adapter denies the request, and the routine is finished. If
12 the HOST_ID of the host controller port requesting access is equal to the HOST_ID in
13 the indexed entry of the virtual port host table, then execution continues from step 326 to
14 step 327.

15
16 In step 327, the microprocessor checks whether the port adapter's physical port
17 "A" or "B" from which the request was received is the same as the physical port "A" or
18 "B" indicated by the PORT A/B flag in the indexed entry of the virtual port host table. If
19 not, then execution branches from step 327 to step 328. In step 328, execution branches
20 from step 328 to step 323 to deny the request if a "fixed port option" is selected. The
21 "fixed port option" would be selected if it is desired for a host controller port to access
22 the specified virtual port only through a specified one of the physical "A" or "B" ports of

1 the port adapter. This option could apply or not apply to all of the host controller ports,
2 or the option could be specified by a flag (not shown) in each entry of the virtual port host
3 table to selectively apply the option to each host controller port. For example, the option
4 could be selected for desktop workstations, and not be selected for portable host
5 computers. If the fixed port option is not selected, then execution continues from step
6 328 to step 329. In step 329, the microprocessor switches the PORT A/B flag in the table
7 entry. Execution continues from step 329 to step 330 of FIG. 27. Execution also
8 continues from step 327 to step 330 of FIG. 27 if the request was received from the port
9 adapter's physical "A" or "B" port indicated by the PORT A/B flag in the indexed entry
10 of the virtual port host table.

11

12 Referring to FIG. 28, in step 330 the S_ID from the request is compared to the
13 S_ID in the indexed entry of the virtual port host table. If the S_ID of the request does
14 not match the S_ID in the indexed entry, then execution branches from step 330 to step
15 331 to recover from the unreported state change. If the port adapter cannot verify the new
16 S_ID in the request, then the request is denied; otherwise, the S_ID value in the indexed
17 entry is changed to the new value in the request, and execution continues in step 332.
18 Execution also continues from step 330 to step 332 if the S_ID from the request matches
19 the S_ID in the indexed entry.

20

21 In step 332 execution branches to step 333 for processing a "Report LUNs"
22 request from a host controller port. In step 333, the port adapter reports to the host

1 controller port the LUNs in the LUN TO LOGICAL VOLUME MAP of the indexed
2 entry of the virtual port mapping table, and the routine is finished. Execution branches
3 from step 332 to step 334 for processing a storage access request from a host controller
4 port. In step 334, the microprocessor searches the LUN TO LOGICAL VOLUME MAP
5 of the indexed entry of the virtual port mapping table for the specified LUN to access. If
6 the LUN to access is not found, then execution branches from step 335 to step 336. In
7 step 336, the port adapter reports to the host controller port that the specified LUN to
8 access is unknown, and the routine is finished.

9

10 Execution continues from step 335 to step 337 if the LUN to access is found in
11 the LUN TO LOGICAL VOLUME MAP of the indexed entry of the virtual port mapping
12 table. In step 337, the private/shared flag is inspected for the indexed entry of the volume
13 access table. If the private/shared flag is set, then execution continues from step 337 to
14 step 338. In step 338, the port adapter accesses the logical volume mapped to the LUN
15 specified by the host controller, and the routine is finished. If the private/shared flag is
16 not set, then execution branches from step 337 to step 339 to access locking information
17 in the cache memory for the logical volume mapped to the LUN specified by the host
18 controller. If the volume is private to the host controller, then access is permitted, and
19 execution branches from step 340 to step 348 to access the volume. If the volume to
20 access is public and is already locked in a fashion incompatible with the access requested
21 by the host (e.g., the volume is already write locked and the host controller requests a
22 read or a write access, or the volume is already read locked and the host controller

1 requests a write access) then access not presently permitted. The host controller's S_ID
2 is placed on a wait list, in order to notify the host controller when the logical volume
3 becomes available; execution branches from step 340 to step 341 to temporarily deny
4 access to the volume, and the routine is finished. If the volume to access is public and is
5 not locked or is locked in a fashion compatible with the requested access by the host
6 controller (e.g., the volume is already read locked and the host controller requests a read
7 access), then a lock is granted to the host controller, and execution branches from step
8 340 to step 338 to access the volume, and the routine is finished.

9

10 GRAPHICAL USER INTERFACE FOR VIRTUAL PORTS

11

12 A conventional graphical user interface (GUI) for a cached storage subsystem of
13 the type shown in FIG. 1 includes a grid of logical volumes to storage adapter ports. At
14 each intersection of the grid, the target/LUN is assigned. This provides a mechanism for
15 the system administrator to set and view the mapping of LUNs to logical storage volumes
16 and the storage adapter ports used for accessing the physical storage volumes that make
17 up the logical storage volumes. Referring to FIG. 29, there is shown an example of how
18 the mapping of LUNs to logical storage volumes can be incorporated into a GUI for a
19 cached storage subsystem that uses virtual ports. It is still possible to use a grid for
20 defining the relationship between logical volumes, LUNs, and storage adapter ports, but
21 the GUI display screen is partitioned into groups of V_PORTS by physical port of the
22 port adapters. In this example, each intersection allows the definition of the LUN only.

1 FIG. 29, for example, shows a single grid 345 for the logical volumes accessible by a host
2 controller port linked to the physical port A of a port adapter 1 addressing LUNs at a
3 virtual port 1. A similar grid could be displayed for each virtual port defined in the
4 storage subsystem.

5
6 Referring to FIG. 30, there is shown a GUI display screen for permitting the
7 system administrator to set up the relationship between logical storage volumes of the
8 storage subsystem (the volume source) and the volumes addressed by a host (the volume
9 user). In this example, the GUI display screen includes, on the left, a list 346 in outline
10 form of storage subsystem components down to a set of logical volumes for one virtual
11 port, and on the right, a list 347 in outline form of host components down to a set of
12 LUNs as addressed from one host controller port. The “bus” in the right column refers to
13 a host controller, and the “target” in the right column refers to a port of the host
14 controller. The display starts out with a list of storage subsystems and a list of hosts
15 linked by the data network. The system administrator selects a storage subsystem or host
16 by clicking on it with a pointing device such as a mouse interfaced to the service
17 processor, and expand the list to the next lower level by double clicking on a selected
18 item of the list. By expanding the lists down to the port adapter and host controller port
19 levels, the system administrator clicks on a particular physical adapter port in the left
20 column to get to virtual ports and logical volumes for the particular adapter port, and then
21 clicks on a particular controller bus and target to get the LUNs as addressed by the host.
22 The system administrator selects an unallocated virtual port in the left column for a

1 physical adapter port and then selects a host controller port in the right column to allocate
2 the host controller port to the virtual port. The system administrator could then select a
3 logical storage volume, or selected range of volumes, and select a LUN or range of LUNs
4 as addressed by the host controller port, to establish the LUN to logical volume mapping.
5 For example, in FIG. 30, the GUI display has drawn a line 348 between the selected
6 logical storage volume VOL1 and the selected LUN 2 to show that such a mapping has
7 been established. The system administrator could also select an item on either list and
8 “collapse” the list to go back up to the level in the outline above the selected item. For
9 data security, any unallocated logical volume should be re-formatted or erased of any pre-
10 existing data before being allocated.

11

12 The system administrator could select an item on either list and select a
13 “properties” option to display properties of the selected item. For example, the grid 345
14 of FIG. 29 could be one of the properties of the virtual port 1. The properties of a logical
15 storage volume would include the host controller ports having access to the logical
16 storage volume, and the paths or adapter ports and virtual port through which the host
17 controller ports can access the logical storage volume.

18

19 The preferred embodiment of the graphical user interface used by the system
20 administrator has the following additional functions. For security, password protection is
21 desirable on a user and/or host level. For flexibility, it is desirable to have be able to
22 modify or view a host's volume configuration from local or remote hosts. The systems

1 administrator should have the ability to read and write the physical and virtual port names
2 of storage subsystem as defined in the a configuration database. The configuration
3 database should contain flags that indicate if a volume can be shared. If sharing of a
4 volume is forbidden, the graphical user interface does not allow it to be allocated to more
5 than one host. If it can be shared, the system administrator is notified of this upon
6 configuration creations of shared volumes.

7

8 The graphical user interface may recognize an "Install" command to be use when
9 a new host controller port is introduced. The install command writes the host's
10 configuration information into the volume configuration database.

11

12 The graphical user interface may recognize a "Replace" command to replace a
13 host controller port. The graphical user interface maintains a list of the WWNs for each
14 controller port on the host. The graphical user interface also keeps a historical database
15 of all WWNs ever known to be networked to the storage subsystem Using this
16 information, the graphical user interface can determine that a known host controller port
17 is not bound to a configuration and also a configuration that contains a new host
18 controller port that did not previously exist on the host. As described above with
19 reference to FIG. 18, the system administrator should be involved in the introduction of
20 any new host controller ports into the configuration.

21

22 The graphical user interface may recognize a "Configuration" command to enter

1 or modify a host's volume configuration information for a local or remote host. This
2 configuration information should be stored in the configuration database in the storage
3 subsystem.

4
5 The graphical user interface should recognize a "remove from service" command
6 for use when removing a host or host controller from service. In response to this
7 command, the graphical user interface deletes the associated volume configuration
8 information from the database.

9
10 The graphical user interface may recognize a "View Configuration" command,
11 and respond by displaying a schematic diagram of the fabric or loop in the data network,
12 and displaying attribute information (such as port WWN) according to its availability.

13
14 The graphical user interface may provide separate displays of the assignments of
15 the virtual ports to the host channel ports, the mappings of LUNs to logical volumes, and
16 the allocation of the logical volumes to the storage adapter ports and to the storage
17 devices.

18
19 One the system administrator defines the virtual ports and mappings between the
20 logical volumes and LUNs, the storage subsystem can automatically assign network
21 addresses, WWNs and S_IDs to the virtual ports, and obtain the WWNs and S_IDs of the
22 host controller ports when the host controller ports log in to the network. This could be

1 done in a fashion similar to that described above with respect to FIGs. 16 to 18, by setting
2 unknown host controller port WWNs and S_IDs to null values.

3

4 Since a cached storage subsystem of the type shown in FIG. 1 typically has a
5 small number of physical adapter ports, such as 16, the method of virtual ports can be
6 introduced in such as system using a small number of V_PORTS per physical ports. For
7 example using only 4 V_PORTS per physical port will allow the cached data storage
8 system to 64 ports. Once users would become familiar with the concept of the virtual
9 ports, the number of virtual ports permitted per physical port could be increased, and
10 additional facilities could be introduced for volume configuration and site management of
11 the Fibre Channel network. These facilities could permit involvement of the host in
12 allocating and de-allocating logical storage volumes and changing the configuration and
13 mapping of LUNs to logical storage volumes.

14

15 HOST INVOLVEMENT IN VOLUME CONFIGURATION AND MAPPING.

16

17 A number of facilities could be used to permit host involvement in the storage
18 volume configuration and mapping. The facilities should not assume the existence of
19 “trusted” hosts. The identity of a host requesting a change in the volume configuration
20 and mapping should be authenticated.

21

22 A primary copy of the configuration information for the volumes accessible to a

1 host is kept in the storage subsystem and on the host. The host should be able to access
2 the primary copy on the storage subsystem if a host's local copy is not available.
3 However, an exception to this would be a host that is responsible for storing the primary
4 copy of another host. The configuration information should be accessed through read and
5 write commands and should be protected against accidental writes. For example, the
6 configuration information is stored in a predefined logical volume, such as a volume
7 accessed at LUN0, that functions as a gatekeeper device.

8
9 No host should be able to gain access to another host's volume unless sharing is
10 enabled. Sharing should be permitted when desired. For example, the volume attributes
11 and locking information (89 in FIG. 7) associated with the logical volumes includes a flag
12 for each volume to indicate whether or not sharing of the volume is permitted. This flag
13 is inspected when the graphical user displays the logical volumes available for allocation,
14 or when a host requests allocation of a volume. A host controller driver program also
15 may limit (or filter) the number of volumes that can be seen by the host operating system
16 (and thus the host applications) by using the information contained in the configuration
17 database.

18
19 Each host should be allowed to use its own volume address space (LUNs)
20 regardless of the volume address space used in the storage subsystem. This is the
21 function of mapping LUNs to logical volumes, as described above.

22

1 Management of the configuration information should be permitted from any
2 location within the data network enterprise. A system administrator should not be
3 required to log into any host that is the target of the configuration changes.

4

5 The facilities should work well with a variety of existing host and storage system
6 components, for example, by preserving SCSI conventions, and supporting both IP and
7 SCSI conventions within a single host controller.

8

9 The facilities should provide for installation, volume reconfiguration, topology
10 reconfiguration, boot from local disk, boot from Fibre Channel disk, normal I/O
11 operation, replacement of a controller from existing host, replacement of a controller
12 known in the current configuration with an unknown host during installation, and reuse of
13 a controller known in the current configuration as a replacement controller. Suitable
14 installation, volume configuration, normal I/O operation, and controller replacement
15 facilities have been described above.

16

17 Topology reconfiguration occurs whenever a connection is added, deleted, or
18 modified in the data network. Some of the types of problems that can occur is for the
19 S_ID of a host controller port to change. Therefore, the ports should login to the data
20 storage subsystem to register the new S_IDs. Otherwise, the port adapters may recognize
21 and recover from unreported state changes, as described above. A host controller port
22 may become unlinked from a physical "A" port and relinked to the physical "B" port. If

1 the “portable host” option is selected, this causes no difficulty. A host controller port
2 may become unlinked from the physical port of a first port adapter and relinked to the
3 physical port of a second port adapter. Unless the same virtual port is defined as
4 accessible to the host controller port in each of the port adapters, this will cause the
5 second port adapter to reject host requests. In this case, it is necessary to transfer the
6 virtual port definitions and mappings from the first storage adapter to the second storage
7 adapter.

8
9 The facilities that can be used by a host for booting are the routine described
10 above for reporting to a host the virtual ports available to the host, and the routine
11 described above for reporting to the host the LUNs available to the host from a virtual
12 port. These routines could be programmed into each port adapter. Alternatively, a host
13 could read the primary copy of the configuration information in the “gatekeeper” volume
14 of storage in the storage subsystem. In any case, the host must be programmed to seek
15 out the LUNs that it can access. For example, a host controller routine first powers up
16 and logs in to the network. Then a mapping driver in the host is loaded into the host’s
17 memory, and the mapping driver contains the network addresses of the data storage
18 subsystem adapter ports to access, and contains instructions for sending the commands to
19 these adapter ports to obtain the LUN information or read the primary copy of the
20 configuration information in the storage subsystem. The host operating system invokes
21 the mapping driver to obtain the LUNs accessible to it.

22

1 It is also desirable to provide the host with the capability of using a logical
2 volume in the storage subsystem as a “boot disk.” Booting over Fibre Channel allows
3 support of diskless systems and, more importantly, provides reliability and a means for
4 disaster recovery. Booting over the Fibre Channel is especially useful when the host is a
5 commodity server. In this situation, the storage subsystem can provide the disk storage
6 for a plurality of the commodity servers. The commodity servers can handle wide area
7 network communication protocols with other hosts such as workstations, and file system
8 management for file systems stored in the logical volumes of the storage subsystem.

9
10 The use of a logical volume in the storage subsystem as a “boot disk” increases
11 reliability of a host because the standalone disk in a typical server represents the
12 component with the worst mean time before failure. In the past, the system administrator
13 has been willing to treat a local disk failure as a server failure. The storage subsystem,
14 however, has a fault-tolerant architecture and uses techniques such as mirroring of the
15 storage devices or RAID so that a logical volume in the storage subsystem has a much
16 higher mean time before failure than the typical server or server disk. The storage
17 subsystem also enhances disaster recovery capabilities because it is easy to recover from
18 the destruction of a commodity server simply by plugging in a replacement commodity
19 server having the same hardware configuration.

20
21 A host computer boots from a present disk address. This is accomplished through
22 one of the following methods. In the PC environment, the host controller is a SCSI

1 controller that must act like an IDE controller and answer all interrupts that are targeted
2 for the disk. The host controller uses a specific SCSI target that has either been
3 hardwired in the design of the host controller stored in nonvolatile memory on the host
4 controller. In other environments, the motherboard of the host typically contains the
5 SCSI address to use in non-volatile memory. The address must be available before any
6 information on a disk is available to be read. This means that the nonvolatile memory
7 must be used either on the motherboard or the host controller.

8
9 The address of the boot disk in the Fibre Channel environment is a storage
10 subsystem port WWN and a storage subsystem LUN. The nonvolatile memory must
11 store a complete path to the boot volume, including the SCSI address to be used by the
12 host as well as the storage subsystem WWN and LUN.

13
14 To support booting from a storage subsystem volume, the host controller is
15 provided with a facility to read and write the boot information. This information should
16 also be saved in the storage subsystem volume configuration database in order that a
17 replacement host controller can be loaded with the correct information during
18 configuration. To facilitate this, the read and write boot commands would each use the
19 following structure:

20
21 struct (int64 hba_wwn; /* which host controller port to address with this
22 cmd */

```

1          int bus;          /* the value for the host bus */
2          int target;       /* the value for the host target */
3          int lun;          /* the value for the host lun */
4          int64 symm_wwn;    /* the value for the storage subsystem port WWN */
5          int symm_lun;      /* the LUN to use on the storage subsystem port */

```

6

7 HOST REQUESTS FOR DYNAMIC ALLOCATION AND DE-ALLOCATION 8 OF LOGICAL VOLUMES

9

10 After an initial configuration by the system administrator and during normal
11 operation of the data processing system, a host application may need additional storage
12 volumes, or may no longer need storage volumes allocated to it. To avoid significant
13 involvement of the system administrator in this situation, the port adapters of the storage
14 subsystem can be programmed to recognize “mount” and “unmount” commands from the
15 host controller. These commands can be similar in form and function to the Unix
16 mount/unmount commands, and would be sent to the storage subsystem gatekeeper (e.g.,
17 at LUN0). A command line for the mount command, for example, has the form:

18

```

19          storage_subsystem_mount volume_group_name host_controller host_lun1 ...
20  host_lunN

```

21

22 A command line for the unmount command, for example, has the form:

1

2 Host controller ID (the value assigned by the switch at login time)

3

4 Hostname IP address

5

6 Volume group name

7

8 Host LUN 1

9

10 ...

11

12 Host LUNn

13

14 In the storage subsystem, the gatekeeper facility responds to the mount command
15 by allocating free logical storage volumes to the specified LUNs, and creating an entry in
16 the volume access table or tables for the specified volume group name and the LUN to
17 logical volume mappings. The entry would include the host controller port's S_ID and
18 WWN. The host could specify attributes for logical volume to be allocated to each
19 volume, such sharing enabled, or sharing enabled for read-only access. In this example,
20 the volumes allocated to the host would all be private, and the private/shared flag for the
21 volume group name would be set. The system administrator would set up in advance any
22 assignment of already volumes to be shared between different host controller ports. It

1 would be possible, however, to have an option whereby one host controller port could
2 request and obtain mounting of volumes shared with another host controller port. If LUN
3 to logical volume mapping is used, this option could be limited to sharing between two
4 ports of the same host controller, or two ports of the same host, because such an option
5 assumes that the LUN to logical volume mapping would be the same for the two host
6 controller ports. In response to such a command, the gatekeeper would check whether the
7 attributes specified in the command for the LUNs to be assigned (such as shared, shared
8 for read-only access) match those of the already allocated LUNs.

9
10 The mount command could also have an option to define the volume group name
11 as local to just the port adapter that receives the command, or global for all of the port
12 adapters in the storage subsystem. In the local case, an entry for the volume group name
13 would be created only in the volume access table stored in the port adapter that received
14 the mount command. In the global case, an entry would be created for the group name in
15 the volume access table in each port adapter of the storage subsystem. In any case, it
16 would be desirable for a host to or host controller have the ability to access the volumes
17 assigned to it through more than two of the ports in order to provide load balancing. This
18 would include not only resending a data access request to another storage subsystem
19 adapter port in response to a port adapter busy signal, but also spreading the host's load
20 over a multiplicity or all of the storage subsystem adapter ports, for example in a round-
21 robin fashion, to decrease the likelihood of any adapter port becoming busy. Such a
22 method can be employed by a host to provide load balancing without any knowledge of

1 the actual loading conditions on any of the ports, and without any coordination with the
2 storage subsystem or other hosts. The round robin technique can be implemented simply
3 by the host or host controller incrementing an index prior to sending each data access
4 request, testing whether the index is equal to the number of subsystem adapter ports
5 accessible to the host and if so resetting the index, and then directing the host's data
6 access request to the accessible storage subsystem adapter port specified by the index.

7

8 The gatekeeper facility responds to the unmount command by checking whether
9 the specified LUNs are all the LUNs assigned to the volume group name in the volume
10 access table, and if so, removing the volume group name, and deallocating the volumes
11 allocated to the group name. Normally, private volumes are erased when they are
12 deallocated. There also could be an option to deallocate some but not all of the logical
13 volumes assigned to a volume group name, and in this case only the entry for the volume
14 group name would remain, but the volume list or LUN to logical volume map for the
15 entry would change.

16

17 When a controller is replaced, the "loss of light" condition (when the controller is
18 taken off line) will cause the automatic logout of the port from the switch. (If a fabric
19 exists, the problem of propagating the information through a fabric switch has not yet
20 been solved by the FC standard.) This information will be passed on to the storage
21 subsystem, which will remove the entry from the table. After the host is rebooted, the
22 mount command will be executed and the connection will be re-established.

1

2 Some hosts such as servers should be continuously logged in to the storage
3 subsystem in the absence of a failure condition or planned removal from the data
4 processing system. In this case, the host should unmount all of their assigned volumes
5 before logging out for a planned removal. To respond to a failure condition, the storage
6 subsystem could be programmed to respond to a report of a state change indicating that
7 such a host has been disconnected from the data network by checking the volume access
8 table or tables for any volume group names corresponding to the host, and if any such
9 volume group name has been found, reporting the error to the system administrator. The
10 system administrator can then decide whether any volumes allocated to the host should be
11 deallocated and erased for use by other hosts.

12

13 An example of a procedure using the host and storage subsystem facilities during
14 a first-time host installation in the data network is as follows:

15

16 1. The host controller of the host powers up. It performs its own initialization
17 but does not execute any operations related to volume configuration.

18

19 2. The host boots over a host CD-ROM drive or local disk (for example in the
20 case of a Fibre Channel storage network added to an existing system) or network storage
21 such as a logical volume in the storage subsystem.

22

- 1 3. Mount a network disk containing the host facilities (seen by the host as an
2 application program) or load the host facilities from the host CD-ROM drive.
3
- 4 4. Start a graphical user interface for the host facilities and enter an Installation
5 Mode.
6
- 7 5. The host installation facility of the graphical user interface queries each host
8 controller port for its port WWN.
9
- 10 6. The host installation facility interrogates the data network for the WWNs of
11 the storage subsystem ports. If the host controller port is in a Fibre Channel loop, then
12 host installation facility probes the loop through a host controller for the WWNs...
13 Otherwise, the host installation facility calls the fabric name server to obtain a list of list
14 of SCSI device WWNs, and the full list is searched for WWNs storage subsystem ports.
15
- 16 7. The host installation facility reads the storage subsystem port text names, for
17 example from block zero of the port configuration data in each storage subsystem's
18 "gatekeeper" volume at LUN0. If the port text name entries are zero, indicating that they
19 have never been entered by the system administrator, then a default of
20 <storage_sybsytem_name>/<storage_subsystem_adapter_port #> is used.
21
- 22 8. The host installation facility displays to the host user each storage subsystem

1 adapter port WWN with its corresponding text name.

2

3 9. When the host user selects a given storage subsystem adapter port WWN, the
4 host installation facility executes a “Report LUNs” command. Using an application
5 interface (API) function called getallprtlun(), all of the LUN bitmap entries and their
6 corresponding volume group names are read from the storage subsystem volume
7 configuration database in the storage subsystem’s “gatekeeper” volume at LUN0. The
8 logical OR of all of the bitmasks from each entry provide a single bitmask of all allocated
9 volumes. While reading each entry, the host installation facility also checks each
10 hostname as it appears in each volume group name returned from the API. This is
11 compared with the current hostname in order to look for previous entries in the
12 configuration database for this host. Upon completion of this sequence of actions, the
13 host installation facility holds the following information:

14

15 All volumes available on the storage subsystem adapter port;

16

17 All volumes on the storage subsystem adapter port that are already allocated to
18 other hosts;

19

20 All volumes that are already allocated to this host for this storage subsystem port;

21

22 Port WWNs for each host controller port of this host; and

1
2 Port WWNs and text names for all storage subsystem ports available.

3
4 10. The host user can create one set of allocated volumes for each host controller
5 port residing on the host. The host facility graphical user interface associates an index
6 number with each volume set, and displays to the host user each index number alongside
7 each host controller port WWN. The graphical user interface also displays the volumes
8 allocated to the host and those volumes that are not allocated to any host. The graphical
9 user interface also displays, on demand, a list of volumes that are allocated to other hosts.
10 From these, the user can select volumes to be allocated to the a selected host controller
11 port of the host. If the volume is already allocated to another host and is marked as
12 unshareable, an error is generated.

13
14 11 If the host user initiates a save operation, the host installation facility writes all
15 the configuration information along with the host controller port WWN to the storage
16 subsystem volume configuration database through the gatekeeper. The storage subsystem
17 makes corresponding revisions to the volume access tables in its cache memory and port
18 adapter memory.

19
20
21 **HOST AUTHENTICATION**
22

1 In a storage system network, is desirable to authenticate host identity to prevent
2 tampering with especially critical data such as the configuration information, which could
3 enable unauthorized access to the most sensitive volumes of storage in the system.
4 Therefore, whenever a host requests a change in the configuration information (i.e., write
5 access), it is desirable to use a challenge-response protocol by which the host must
6 identify itself in response to a request from the storage subsystem. Preferably this
7 challenge-response protocol incorporates a machine-based security feature that generates
8 a response that cannot be re-used if intercepted.

9

10 The machine-based security feature includes an encryption key stored in memory
11 of the hosts or host controllers. Each host controller, for example, has a unique
12 encryption key that is different from the keys of the other host controllers. Preferably the
13 encryption key is stored in “write only” memory of a monolithic semiconductor
14 integrated circuit chip (hereinafter referred to as an “identity” chip) in such a way that the
15 stored encryption key can only be read by encryption circuitry on the chip. The “write
16 only” memory could be an electrically erasable and programmable read-only memory
17 (EEPROM) array with a metal shielding layer over the memory array so that the stored
18 key would be virtually impossible to recover by probing, inspection, disassembly, or
19 “reverse engineering” of the identity chip. In this case, would be possible to publish the
20 method by which the identity of the chip is authenticated. The identity chip, for example,
21 is a microcontroller having such a write-only EEPROM program memory programmed to
22 perform a unique encoding function corresponding to the key. The “write only” memory

1 and the encryption circuitry could be a small portion of the circuitry on the identity chip.
2 For example, the identity chip could be a microprocessor chip used as the main processor
3 of a host controller circuit card or workstation.

4

5 By incorporating the identity chip into the host controller, host motherboard, or
6 the main processor chip of the host, is possible to verify the identity of the host controller,
7 host motherboard, or main processor chip of the host . The identity chip, however, does
8 not verify the identity of the user of the host. To guard against unauthorized users, it is
9 still necessary to verify the identity of the user. User verification would be done by the
10 operating system in the usual fashion, for example at a user login time when the user is
11 requested to enter a password.

12

13 The identity chip could be attached to or incorporated into any kind of object or
14 thing in order to authenticate the identity of that object or thing. For example, the chip
15 could be embedded in a debit or credit card to authenticate the identity of the card before
16 debiting or charging a corresponding account. The chip could be attached to an object or
17 thing together with a wireless data transceiver to permit remote interrogation of the chip,
18 for example in a tag implanted under the skin of an animal to remotely authenticate the
19 identity of the animal. The identity chip could store descriptive information about the
20 object or thing to which is attached, in order to permit verification that the identity chip
21 has not been transferred to a different object or thing. If the descriptive information is
22 unique to the particular object or thing, then the descriptive information could be

1 incorporated into the secret key. For example, if the identity chip is used in a tag for an
2 animal, part of the key could include a sequence of the genetic code of the animal.

3

4 FIG. 29 shows a preferred construction for the identity chip. The identity chip
5 350 includes encryption circuitry 351 and a "write-only" EEPROM memory 352 for
6 storing at least one key 353. The key 353 can be written into the memory 352 from a data
7 path 354 including external leads connected to the chip. The key 353 can be read from
8 the memory 352 by the encryption circuitry, but the key cannot be read from the data path
9 354 or any other external leads connected to the chip. For example, the encryption
10 circuitry 352 includes a microprocessor 355 and a microcode read-only memory (ROM)
11 356 storing microcode executed by the microprocessor. The microprocessor is
12 programmed to recognize a command from the data path 354 for writing into the memory
13 352 a key from the buys 354. The microprocessor is also programmed to recognize a
14 command from the data path 354 for receiving a number from the data path, reading the
15 key 353 from the memory 352, encrypting the number with the key, and transmitting the
16 encryption result onto the data path. The microprocessor, however, will not recognize
17 any command for transmitting the key onto the data path 354 or any other leads of the
18 chip. In this sense, the memory 352 is a "write-only" memory. Moreover, the EEPROM
19 memory 352 and at least the internal data path 357 to the memory 352 are covered by an
20 upper layer of metal 358 (shown in dashed lines in FIG. 31) on the chip 350 so that it is
21 virtually impossible for the key to be recovered by probing, inspection, disassembly, or
22 "reverse engineering" of the chip. The EEPROM 352 could store a plurality of different

1 keys, and the microprocessor could recognize a command from the data path 354 for
2 selecting which of the keys to use for encrypting the number received on the data path.

3

4 FIG. 32 shows how identity chips 361, 363 are incorporated in the data processing
5 system of FIG. 1 for authenticating host controller identity. In this example, the Fibre
6 Channel loop 41 is used for access to highly sensitive data, such as configuration
7 information, stored in the cached storage subsystem 20. The identity chip 361 in the host
8 controller 61 stores a unique secret key 362 for the host controller 61, and the identity
9 chip 363 in the host controller 62 stores a unique secret key 364 for the host controller 62.
10 Copies of the keys 362, 364 are stored in a list 365 in the host controller. The list 365, for
11 example, is a table, and each key in the list 365 is associated with the WWN of a
12 respective host controller known to the port adapter. A copy of the list 365 of keys is
13 stored in a logical storage volume of the storage subsystem, or the list of keys is stored in
14 a nonvolatile portion of the port adapter memory 77, so that the list will be retained if
15 power to the port adapter is lost. Such a nonvolatile portion of the port adapter memory
16 77 could be provided by one or more identity chips constructed as shown in FIG. 29. The
17 port adapter memory 77 further includes a list 366 of random numbers sent to the hosts.
18 The list 366, for example, is maintained as a queue. The list of random numbers 366 and
19 the list of keys 365 are used by host authentication routines 367 in the port adapter
20 microcode 79. As the random numbers are received by the hosts, the host controllers
21 place the random numbers in similar lists 368, 269 in their respective memories 370, 371.
22 As used in this specification, the term "random number" would include a so-called

1 “pseudo-random number” so long as the number would appear to be selected at random
2 during the typical duration of time that a host controller is logged in to a port adapter.
3 For example, the random numbers transmitted by the port adapter to a host controller are
4 generated by a conventional random number generator routine that is seeded at the time
5 that the host controller logs in to the port adapter. The BASIC programming language,
6 for example, provides such a random number generator routine.

7
8 Shown in FIG. 33 is a flowchart of the challenge-response protocol. The actions
9 of the port adapter are shown on the left side of the flowchart, and related actions of a
10 host controller are shown on the right side of the flowchart. The steps of the left side of
11 the flowchart, for example, are performed by one of the host authentication routines (367
12 in FIG. 32). The challenge-response protocol is used, for example, when a host controller
13 logs in to the port adapter. Prior to login, however, the system administrator must load
14 the host controller keys into the port adapter and host controllers. This should be done in
15 a secure fashion, to guarantee that the keys cannot be intercepted when they are
16 distributed to the host controllers. The system administrator could load the keys into the
17 port adapters through the service processor when assigning the logical storage volumes to
18 the host controllers. In a similar fashion, the keys could be entered into each host
19 controllers through a dedicated I/O port to avoid any need to transmit the keys over the
20 data network.

21

22 In a first step 381, the host controller sends a request to the port adapter for some

1 work. In step 382, the port adapter receives the work request. In step 383, the port
2 adapter responds to the work request by sending back a “challenge” request including a
3 random number that is used as a one-time password. Since the number is random and
4 changes every time it is needed, there is no way to predict it, nor record it and provide a
5 false reply to a request for authentication.

6
7 In step 384, the host controller receives the random number and the request for
8 verification from the port adapter. In step 385, the port adapter encrypts the random
9 number with its copy of the host controller’s key. Concurrently, in step 386, the host
10 controller encrypts the random number with its copy of the encryption key. The
11 encryption process can use any kind of encryption technique since the result never has to
12 be decrypted. For example, the encryption could use the well-known method of the Data
13 Encryption Standard (DES), or any sequence of encryption operations, for example,
14 substitution, blocking, permutation, expansion, and compaction. It is also possible to use
15 the well-known RSA public/private key system. In the conventional use of the RSA
16 system, public keys are used for encryption, and a private keys are used for decryption.
17 Therefore, if the port adapter and the host controller used RSA system, the public keys
18 would be kept secret and would used for encoding, and the private keys would not be
19 used.

20

21 In step 387, the host controller sends the encryption value to the port adapter. In
22 step 388, the port adapter receives the encryption value from the port adapter. In step

1 389, the port adapter compares the encryption value that it computed with the encryption
2 value received from the host controller. If they match, then the host controller is
3 authenticated, and the host controller is granted access to the storage subsystem. If they
4 do not match, the host controller is denied access to the storage subsystem.

5
6 One preferred way of applying the “challenge-response” protocol in a Fibre
7 Channel network is by using the Exchange IDs (Originator and Responder) within the
8 Fibre Channel header of the Frame Structure. The primary function of the Fibre Channel
9 network is to receive Frames from a source port and route them to a destination port. As
10 shown in Fig. 30, a Frame 400 includes information to be transmitted (Payload 401), the
11 source address (S_ID 402), the destination address (D_ID 403), and various kinds of link
12 control information.

13
14 Each Frame begins with a Frame Delimiter 404 indicating the start of the Frame.
15 The Frame Delimiter 404 is immediately followed by a Frame Header 405. The Frame
16 Header is used to control link applications, control device protocol transfers, and detect
17 missing or out of order Frames. The Frame Header includes the source address (S_ID
18 402) and the destination address (D_ID 403). An optional header 406 may contain
19 additional link control information. The Payload 401 follows the Frame Header or the
20 optional header. Four bytes of Cyclic Redundancy Check (CRC) follows the Payload.
21 The CRC is used to detect transmission errors. The Frame 400 concludes with an End of
22 Frame delimiter 408.

1

2 The Exchange_ID 409 is an end portion of the Frame Header 405. The
3 Exchange_ID includes an Originator_ID 410 and a Responder_ID 411. The originator of
4 a communication sends the first message with a value specified as the Originator_ID.
5 The Originator_ID must be unique amongst all the currently active exchanges. The
6 responder receives the header. When it sends back a message, it places the Originator_ID
7 plus its own Responder_ID into the header. This ID pair now comprises a unique
8 Exchange_ID.

9

10 Once a host controller has logged in to a port adapter, it is desirable to
11 authenticate the host controller for each storage access request. If the challenge-response
12 protocol of FIG. 33 were used for each storage access request, however, the traffic on the
13 data network would be substantially increased because the Frame for each storage access
14 request from the host controller would be followed by a Frame for the “challenge”
15 message from the port adapter, and a Frame for the “response” of the host controller. In
16 order to eliminate the “challenge” and “response” Frames, the port adapter and host
17 controllers maintain and use the lists 366, 368, and 369 of random numbers so that a host
18 controller can include an encrypted value in the Exchange_ID of each storage access
19 request.

20

21 Referring to FIGs. 35 and 36, there is shown a flow chart of a preferred procedure
22 for maintaining and using the lists 366, 368, and 369 of the random numbers. In a first

1 step 421 of FIG. 35, a host controller sends a login request to a port adapter. In step 422,
2 the port adapter receives the login request from the host controller. In step 423, the port
3 adapter authenticates the host, for example by using the “challenge-response” protocol of
4 FIG. 33. Then in step 434, the port adapter sends an initial set of random numbers to the
5 host controller. The port adapter places these random numbers in its list 377. The list
6 377 is maintained as a set of first-in, first-out (FIFO) queues including a respective queue
7 allocated for the random numbers sent to each host controller. In a similar fashion, the
8 host controller places the initial set of random numbers it receives in to its list 368 or 369
9 of random numbers. The host controller maintains its list as a set of FIFO queues
10 including a respective queue for each port adapter that it has logged in to. The size of
11 each FIFO queue, for example, is equal to the size of a queue that holds outstanding
12 requests from the adapter ports or host controller ports. This is done so that the port
13 adapter and host controller can encrypt the random numbers in their respective queues as
14 a background task relative to the outstanding requests. The host controller should
15 typically produce the encryption value by the time that the host controller receives a
16 request from the host to transmit a corresponding storage access request to the port
17 adapter, and the port adapter should typically produce an encryption value by the time
18 that the port adapter receives the corresponding encryption value in a storage access
19 request from the host controller. The port adapter begins encryption in step 426, and the
20 host controller begins encryption in step 427.

21

22 Every time the host needs to send an authenticated storage access request to the

1 storage subsystem, as shown in step 428, the host controller takes the encrypted value of
2 the next one of the random numbers in the FIFO queue for the port adapter, and uses the
3 encrypted value as the Originator_ID in the Frame for the storage access request. In step
4 429, the port adapter receives the Frame from the host controller. In step 430, the port
5 adapter compares the encrypted value in the Frame with its encryption of the next random
6 number in its FIFO queue for that particular host controller. If there is no match, then the
7 storage subsystem rejects the message. This is an error condition that is reported back to
8 the host controller. If there is a match, then the requested storage access can proceed, as
9 shown in FIG. 36. In step 431, the port adapter removes the last random number from the
10 FIFO queue for the host controller, and generates a new random number used as the
11 Responder_ID in its reply to the data access request from the host controller, and places
12 this new random number at the end of its list, in its FIFO queue for the host controller. In
13 step 432, the host controller receives the response from the port adapter, and places the
14 new random number at the end of its list, in its FIFO queue for the port adapter. In step
15 433, the port adapter continues encrypting the random numbers in its list, in the order in
16 which the random numbers were put into the list. Concurrently, the host controller
17 continues encrypting the random numbers in its list. In this fashion, the host controller
18 will have in advance a random number to use as the Originator_ID in a next message, and
19 depending on the priority of any other tasks, the host controller can encrypt the random
20 number well before the encryption need be inserted into the next message. This ensures
21 that the authentication method will not significantly impact the storage access time in the
22 data processing system.

STORAGE NETWORK CONFIGURATIONS

FIG. 2 introduced a storage system configuration in which a respective Fibre Channel loop 41, 42, 43, 44 was linked to a storage subsystem adapter port, and a number of hosts 22, 23, 24, 25 were linked to each loop, and each host was linked through a respective controller port to two loops. Such a fault-tolerant network configuration should support from two (4K I/O) to ten (32K I/O) small server hosts per storage subsystem adapter port. The storage subsystem 20, for example, can be fully loaded with sixteen adapter ports, which should support from 32 to 160 small server hosts.

In contrast to small server hosts, workstation hosts have a relatively light loading on the storage subsystem adapter ports. Also, fault tolerant links usually are needed between the workstations and the storage subsystem adapter ports. Therefore, as shown in FIG. 37, it is possible to link a large number of workstations 501, 502, 503, 504 to a cached storage subsystem 505 through a respective loop 506, 507 linked to each adapter port. For example, in the storage network configuration of FIG. 37, each of sixteen loops could be linked to about fifty workstations, to network a total of about eight-hundred workstations to one storage subsystem.

In a file-only environment (i.e., no video service), a single storage subsystem is capable of supporting considerably more than eight-hundred workstations. If it is desired

1 to network more than eight-hundred workstations to one storage subsystem, then a
2 maximum of about fifty workstations per loop becomes a constraint. The number of
3 workstations can be increased while meeting this constraint by using switches to increase
4 the number of loops over the number of adapter ports. As shown in FIG. 38, for example,
5 three sixteen- port switches 511, 512, 513 expand the number of ports connectable to
6 loops from the sixteen of the cached storage subsystem 514 to fifty-five. Each of the
7 three sixteen port switches 511, 512, 513 has three of its ports connected directly to three
8 respective storage subsystem adapter ports, and thirteen of its ports connected directly to
9 loops, such as the loops 515, 516, 517 shown directly connected to a port of a respective
10 one of the switches 511, 512, 513. This would permit the fifty-five loops to support a
11 total of thirty-three hundred workstations, although the loops would not have similar
12 performance since thirty-nine of the loops would be linked directly to ports of the
13 switches 511, 512, 513, and seven of the loops would be linked directly to storage
14 subsystem adapter ports.

15

16 Another environment where multiple hosts are connected to a storage subsystem
17 is known as the cluster of storage area networks (SAN) environment. In this case,
18 multiple hosts are linked to more than one storage subsystem, and the storage subsystems
19 are linked directly for the transfer of remote dual-copy data. This typically is not a high
20 performance environment, but fault tolerance is a requirement. Therefore, there is a need
21 for more than one path from each host to each storage subsystem. Shown in FIG. 39, for
22 example, is a two-loop storage network configuration including a first cached storage

1 subsystem 531, a second cached storage subsystem 532. The two cached storage
2 subsystems are coupled by a direct link 533 for transfer of remote dual-copy data. The
3 network configuration includes sixteen hosts, only the first host 534 and the sixteenth
4 host 535 being shown. A first loop 536 has links to one port on each of the cached
5 storage subsystems 531, 532 and links to one port on each of the sixteen hosts 534, 535.
6 A second loop 537 also has links to another port on each of the cached storage
7 subsystems 531, 532 and to another port on each of the sixteen hosts 534, 535.

8
9 If the network configuration of FIG. 39 would not provide sufficient throughput
10 performance, then additional loops are added, and the number of hosts directly linked to
11 each loop is decreased. Shown in FIG. 40, for example, is a four-loop configuration,
12 including as before two cached storage subsystems 541, 542 having a direct link 543
13 between them, and sixteen hosts, only the first host 544, the eighth host 545, the ninth
14 host 546, and the sixteenth host 547 being shown. Two loops 548, 549 each have one
15 direct link to one respective port on each of the cached storage subsystems 542, 543, and
16 each of the first eight of the sixteen hosts, including, as shown, the first host 544 and the
17 eighth host 545. None of the last eight of the sixteen hosts, including, as shown, the ninth
18 host 546 and the sixteenth host, 547, have a port directly linked to the two loops 548,
19 549. Two loops 550, 551 each have one direct link to one respective port on each of the
20 cached storage subsystems 542, 543 and each of the last eight of the sixteen hosts,
21 including, as shown, hosts 9 to 16. None of the first eight of the sixteen hosts, including,
22 as shown, the first host 544 and the eighth host 545 have a port directly linked to the

1 loops 550, 551.

2

3 CONCLUSIONS

4

5 From the above, it is seen that a number of facilities have described to provide
6 various features and advantages in a data processing system having a large number of
7 hosts networked to one or more cached storage subsystems, some or all of which could be
8 desired or required by any particular system users and managers. It should be apparent to
9 a person of ordinary skill in the data processing art that the embodiments described above
10 can be constructed to meet some or all of the following requirements:

11

12 A. General Requirements

13

14 1. Any number of storage subsystems should be simply plugged into the
15 network and easily configured by a system administrator for host access to selected
16 volumes.

17

18 2. Each entity in the data processing system (storage subsystem, data network,
19 and host) should have separate and distinct management interfaces. The system
20 administrator graphical user interface, or host graphical user interface, may provide a
21 point of integration for the management interfaces of the entities.

22

1 B. Data Network Requirements:

2

3 1. No path management internal to the data network should be required. Path

4 analysis is undesirable for providing connectivity to a host while retaining required

5 performance, in a system having a very large number of hosts, and current network

6 technology such as Fibre Channel fabric switches having at most a few tens of ports. The

7 internal structure of the data network can be complex, and in this case it is relatively

8 difficult to analyze each path inside the data network , from the host to the storage

9 subsystem, to determine the correct port and path for each host. Since the number of

10 hosts is very large, it is also an almost impossible task to directly control the performance

11 requirements for a single host. With a large number of hosts, the performance typically

12 will not cluster, but will have random hot spots constantly changing over time.

13

14 2. Invisible internal path redundancy. There should always be at least two full

15 redundant paths to reach the storage subsystem volumes from an external port of the data

16 network.

17

18 3. The ports and protocol to/from the storage subsystem and data network

19 should comply with current standards, including both fabric and loop. Given the cost of a

20 switch port and the number of hosts involved, it is expected that the fabric will be made

21 up of a combination of loops and fabric direct connect. Each port going to/from the

22 storage subsystem should simultaneously support the different types of protocols. The

1 external port type should be determined by the specific solution.

2

3 4. The storage subsystem should be allowed to register for notification of login
4 state changes. The storage subsystem should keep current state information about the
5 hosts and connections, especially to allow single use of a group of volumes. One
6 limitation of the Fibre Channel protocol is that it does not support a logout. Along with
7 that is the issue of detecting the loss of a host (this functionality is currently being
8 developed in the FC standard).

9

10 5. On-line “hot” management of data network configuration.

11

12 6. Replacement of a host controller or storage subsystem director should not
13 require any manual host-to-volume reconfiguration of the data network:

14

15 7. The switches used in the data network should be interoperable with multiple
16 fabric vendors.

17

18 C. Storage Subsystem Requirements:

19

20 1. It is desirable to the storage subsystem to permit all volumes to be seen
21 through a single storage subsystem port. This eliminates the difficult management
22 problem of configuring and maintaining the storage subsystem and the data network to

1 guarantee a path from a given external port of the data network to a given volume.

2

3 2. It is desirable to have a partitioning feature that will define the set of volumes
4 that can be seen by a single host. This will restrict other hosts from seeing volumes they
5 are not configured to see.

6

7 3. A change of a port ID (source id or worldwide name) in the host should not
8 affect the partitioning definition in requirement 2 above.

9

10 4. It is desirable to have a mapping feature that will allow the host to specify its
11 own LUN, which would be mapped to a logical LUN within the storage subsystem.

12

13 5. It is desirable to have the capability of selecting between “simultaneous multi
14 host access” to a volume, and a “single host access” at a time. This is to control data
15 sharing capabilities.

16

17 6. It is desirable to limit host-to-volume access to read-only or read/write.

18

19 7. Each port should simultaneously support all SCSI address modes.

20

21 8. Each port should simultaneously support both Fibre Channel class 2 and class

22 3.

1

2

9. Both loop and fabric Fibre Channel topologies should be supported.

3

4

5

10. It should be possible to back up and restore the host-to-volume connectivity configuration information.

6

7

8

11. The storage subsystem should allow a volume to be used as boot disk for a host.

9

10

11

12. Replacement of a host controller or storage subsystem director should not require any manual host-to-volume reconfiguration of the storage subsystem.

12

13

D. Host Requirements:

14

15

16

1. The host should be responsible for establishing a connection with the storage subsystem.

17

18

19

2. The host should have the capability of specifying the target LUN information to be used in mapping LUNs to logical storage volumes.

20

21

22

3. The host should provide security for the configuration information (volume group and LUNs) that it uses to establish connectivity with the storage subsystem.